

COMPUTING THE KREISS CONSTANT OF A MATRIX*

TIM MITCHELL†

Abstract. We establish the first globally convergent algorithms for computing the Kreiss constant of a matrix to arbitrary accuracy. We propose three different iterations for continuous-time Kreiss constants and analogues for discrete-time Kreiss constants. With standard eigensolvers, the methods do $\mathcal{O}(n^6)$ work, but we show how this theoretical work complexity can be lowered to $\mathcal{O}(n^4)$ on average and $\mathcal{O}(n^5)$ in the worst case via divide-and-conquer variants. Finally, locally optimal Kreiss constant approximations can be efficiently obtained for large-scale matrices via optimization.

Key words. discontinuity of Kreiss constants, inverses of Kronecker sums, distance to uncontrollability algorithms, transient growth, pseudospectra

AMS subject classifications. 15A16, 37C75, 39A22, 39A30, 65F30, 65F60

DOI. 10.1137/19M1275127

Notation. $\|\cdot\|$ denotes the spectral norm, $\sigma_{\min}(\cdot)$ the smallest singular value, and $\Lambda(\cdot)$ the spectrum. A matrix $A \in \mathbb{C}^{2n \times 2n}$ is Hamiltonian if $(JA)^* = JA$ and symplectic if $A^*JA = J$, where $J = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}$. A matrix pencil $A - \lambda B$, where $A, B \in \mathbb{C}^{2n \times 2n}$, is symplectic if $A^*JA = B^*JB$. Euler’s number, $2.71828\dots$, is denoted by e . I_n is the $n \times n$ identity, though we will often omit the subscript when the dimension is clear.

1. Introduction. Given a matrix $A \in \mathbb{C}^{n \times n}$, the ordinary difference equation

$$(1.1) \quad x_{k+1} = Ax_k$$

is asymptotically stable if A is Schur stable, i.e., if $\rho(A) < 1$, where ρ denotes the spectral radius. While $\rho(A)$ tells one about the asymptotic behavior of (1.1), it does not convey information about its transient behavior. For that, we can look at the Kreiss Matrix Theorem, which says for any matrix $A \in \mathbb{C}^{n \times n}$ [30, eq. 18.2]

$$(1.2) \quad \mathcal{K}(A) \leq \sup_{k \geq 0} \|A^k\| \leq en\mathcal{K}(A),$$

where the *Kreiss constant* $\mathcal{K}(A)$ is given by [30, p. 143]

$$(1.3) \quad \mathcal{K}(A) = \sup_{z \in \mathbb{C}, |z| > 1} (|z| - 1) \|(zI - A)^{-1}\|.$$

As also noted in [30, p. 143], $\mathcal{K}(A)$ has an equivalent formulation of

$$(1.4) \quad \mathcal{K}(A) = \sup_{\varepsilon > 0} \frac{\rho_\varepsilon(A) - 1}{\varepsilon},$$

where the ε -pseudospectral radius ρ_ε is defined by

$$(1.5a) \quad \rho_\varepsilon(A) = \max\{|z| : z \in \Lambda(A + \Delta), \|\Delta\| \leq \varepsilon\}$$

$$(1.5b) \quad = \max\{|z| : z \in \mathbb{C}, \|(zI - A)^{-1}\| \geq \varepsilon^{-1}\}.$$

*Received by the editors July 16, 2019; accepted for publication (in revised form) by M. A. Freitag September 8, 2020; published electronically December 15, 2020.

<https://doi.org/10.1137/19M1275127>

Funding: The author’s visits to the Courant Institute, NYU, where some of this work was conducted, were supported by the National Science Foundation grant DMS-1620083.

†Computational Methods in Systems and Control Theory, Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, 39106 Germany (mitchell@mpi-magdeburg.mpg.de).

For any matrix $A \in \mathbb{C}^{n \times n}$, $\mathcal{K}(A) \geq 1$, and $\mathcal{K}(A)$ may be arbitrarily large. As is well known, a matrix A is *power-bounded*, i.e., $\mathcal{K}(A) < \infty$, if and only if $\rho(A) \leq 1$ and all eigenvalues of A with modulus 1 are nondefective. If A is normal and $\rho(A) \leq 1$, then $\mathcal{K}(A) = 1$.

As discussed in [30, p. 177], the original statement by Kreiss in 1962 [19] actually had a far looser upper bound than (1.2): approximately $c^n \mathcal{K}(A)$. The reduction of the constant factor to its current form in fact occurred over nearly 30 years in at least nine separate steps, with Spijker proving the conjecture of [24, p. 590] to finally obtain the (in a certain sense) tight factor of en in 1991 [29].

The Kreiss Matrix Theorem also comes in a continuous-time variant for an ordinary differential equation

$$(1.6) \quad \dot{x} = Ax,$$

which is asymptotically stable if A is Hurwitz stable, i.e., if $\alpha(A) < 0$, where α denotes the spectral abscissa. In this case, the Kreiss Matrix Theorem states that [30, eq. 18.8]

$$(1.7) \quad \mathcal{K}(A) \leq \sup_{t \geq 0} \|e^{tA}\| \leq en\mathcal{K}(A),$$

where, by [30, eq. 14.7], $\mathcal{K}(A)$ is now equivalently given by either

$$(1.8) \quad \mathcal{K}(A) = \sup_{z \in \mathbb{C}, \operatorname{Re} z > 0} (\operatorname{Re} z) \|(zI - A)^{-1}\|$$

or

$$(1.9) \quad \mathcal{K}(A) = \sup_{\varepsilon > 0} \frac{\alpha_\varepsilon(A)}{\varepsilon},$$

and where the ε -pseudospectral abscissa α_ε is defined by

$$(1.10a) \quad \alpha_\varepsilon(A) = \max\{\operatorname{Re} z : z \in \Lambda(A + \Delta), \|\Delta\| \leq \varepsilon\}$$

$$(1.10b) \quad = \max\{\operatorname{Re} z : z \in \mathbb{C}, \|(zI - A)^{-1}\| \geq \varepsilon^{-1}\}.$$

Like the discrete-time case, $\mathcal{K}(A) \geq 1$ and can be arbitrary large. If A is normal and $\alpha(A) \leq 0$, then $\mathcal{K}(A) = 1$.

Despite the wealth of work done over decades towards making the upper bound of the Kreiss Matrix Theorem now tight, there has been no algorithm given to actually compute $\mathcal{K}(A)$ with guarantees. In the literature, $\mathcal{K}(A)$ is often just approximated by plotting (1.4) or (1.9); e.g., see [13] and [25, Chapter 3.4.1].

In this paper, we propose the first globally convergent algorithms to compute both continuous- and discrete-time Kreiss constants to arbitrary accuracy. We assume that A is nonnormal, as otherwise computing its Kreiss constant just involves checking if A is unstable. Furthermore, we assume that $\alpha(A) < 0$ or $\rho(A) < 1$ holds, respectively, in the continuous- or discrete-time case. With standard eigensolvers, the three different methods we propose all have $\mathcal{O}(n^6)$ work complexities,¹ but by also developing so-called *divide-and-conquer* variants, we show how this theoretical work complexity reduces to $\mathcal{O}(n^4)$ on average and $\mathcal{O}(n^5)$ in the worst case. Our work also shows that

¹We use the standard convention of treating dense computations of singular values, eigenvalues, solutions of Sylvester equations, etc., as *atomic operations* with cubic costs in the dimensions of the associated matrices. We additionally assume these costs become linear in the dimension of the matrices when corresponding sparse methods are available.

locally optimal approximations to $\mathcal{K}(A)$ can be efficiently and reliably obtained for large-scale matrices via standard optimization techniques. Furthermore, we establish some variational properties of Kreiss constants, including that the Kreiss constant of a matrix A is not a continuous function with respect to the entries of A . Finally, as a side effect of our work, we also propose an important modification to the *distance-to-uncontrollability* algorithms of [14, 9, 15] to greatly improve their reliability in practice.

The contributions of the paper are structured as follows. In section 2, we introduce the so-called distance to uncontrollability and present a theorem of Gu [14, Theorem 3.1] and distance-to-uncontrollability algorithms that are based upon it. In section 3, we establish variational properties of Kreiss constants that we will need and show that there is a potentially exploitable similarity to computing the distance to uncontrollability. Then, in section 4, we develop theorems for continuous-time $\mathcal{K}(A)$ that are analogues of the aforementioned theorem of Gu, but nevertheless show that, due to key structural differences, existing distance-to-uncontrollability algorithms will not directly extend to Kreiss constants. By developing a so-called *globality certificate* in section 5, we present our first algorithm for computing continuous-time Kreiss constants, which is an optimization-with-restarts method *using backtracking*, and then establish an asymptotically faster *divide-and-conquer* variant that is inspired by [15]. We then modify the premise of our globality certificate in a crucial way in section 6 to develop a second certificate with significant structural differences and properties; this alternative certificate enables two other algorithms for continuous-time $\mathcal{K}(A)$, an optimization-with-restarts method *without backtracking* and a trisection iteration, which can be considered closer analogues of the distance-to-uncontrollability algorithms of [9]. Faster divide-and-conquer versions of these two methods are also developed. In section 7, we consider the case of discrete-time $\mathcal{K}(A)$ and present discrete-time analogues of all of our continuous-time algorithms and associated theoretical results; to the best of our knowledge, this is also the first extension of the 2D level-set ideas of [14] to a discrete-time setting, and it turns out to have surprising differences. Numerical examples are presented in section 8, with concluding remarks made in section 9.

For this paper, the following general definition and two theorems will be needed. The theorems can be found in several places in various forms, such as [21, 28] and [23, Theorem 13.16], respectively.

DEFINITION 1.1. *Given a domain $\Omega \subseteq \mathbb{R}$, a function $f : \Omega \rightarrow \mathbb{R}$ has a global Lipschitz constant (GLC) of $c \geq 0$ if $|f(x) - f(y)| \leq c|x - y|$ for all $x, y \in \Omega$.*

THEOREM 1.2. *For $x, y \in \mathbb{R}$, let $A(x, y)$ be a twice-differentiable $n \times n$ Hermitian matrix family, and for a point (\hat{x}, \hat{y}) , let $\lambda_1 \geq \dots \geq \lambda_n$ be the eigenvalues of $A(\hat{x}, \hat{y})$ with associated unit-norm eigenvectors q_1, \dots, q_n . Then, assuming λ_j is unique,*

$$\frac{\partial^2}{\partial x \partial y} \lambda_j \Big|_{x=\hat{x}, y=\hat{y}} = q_j^* \frac{\partial^2 A(\hat{x}, \hat{y})}{\partial x \partial y} q_j + 2 \sum_{k \neq j} \frac{q_j^* \frac{\partial A(\hat{x}, \hat{y})}{\partial x} q_k \cdot q_j^* \frac{\partial A(\hat{x}, \hat{y})}{\partial y} q_k}{\lambda_j - \lambda_k}.$$

THEOREM 1.3. *Let $A \in \mathbb{C}^{n \times n}$ with $\lambda_j \in \Lambda(A)$ for $j = 1, \dots, n$ and $B \in \mathbb{C}^{m \times m}$ with $\mu_k \in \Lambda(B)$ for $k = 1, \dots, m$. Then the Kronecker sum $A \oplus B = I_m \otimes A + B \otimes I_n$ has eigenvalues $\lambda_j + \mu_k$ for all pairs of j and k .*

2. Computing the distance to uncontrollability. Given $A \in \mathbb{C}^{n \times n}$ and $B \in \mathbb{C}^{n \times m}$, consider the linear control system

$$(2.1) \quad \dot{x} = Ax + Bu,$$

where the state $x \in \mathbb{C}^n$ and $u \in \mathbb{C}^m$, the *control input*, are both dependent on time. The system (2.1) is *controllable* if given respective initial and final states $x(0)$ and $x(T)$ with $T > 0$, there exists a control $u(\cdot)$ that realizes some trajectory $x(\cdot)$ with endpoints $x(0)$ and $x(T)$. The *distance to uncontrollability*, which we denote as $\tau(A, B)$, can be computed via solving the nonconvex optimization problem [12]

$$(2.2) \quad \tau(A, B) = \min_{z \in \mathbb{C}} \sigma_{\min} \left(\begin{bmatrix} A - zI & B \end{bmatrix} \right) = \min_{x, y \in \mathbb{R}} f(x, y),$$

where $f(x, y) = \sigma_{\min}(F(x, y))$ and $F(x, y) = \begin{bmatrix} A - (x + iy)I & B \end{bmatrix}$. The first practical algorithm to address computing $\tau(A, B)$ is due to Gu [14], based on the following result [14, Theorem 3.1].

THEOREM 2.1. *Let $\gamma, \eta \geq 0$ be given. If $\tau(A, B) \leq \gamma$ and $\eta \in [0, 2(\gamma - \tau(A, B))]$, then there exists a pair $x, y \in \mathbb{R}$ such that*

$$(2.3) \quad f(x, y) = f(x + \eta, y) = \gamma.$$

COROLLARY 2.2. *Let $\gamma, \eta \geq 0$ be given. If there do not exist any pairs $x, y \in \mathbb{R}$ such that (2.3) holds, then*

$$(2.4) \quad \tau(A, B) > \gamma - \frac{\eta}{2}.$$

The proof of Theorem 2.1 relies on the fact that $f(x, y)$ has a GLC of 1 with respect to either x or y .

What [14, section 3.2] additionally devised was a sequence of computations to verify whether either (2.3) or (2.4) holds for a given choice of γ and η . This verification procedure, *when using exact arithmetic*, is able to detect and find any points (x, y) such that (2.3) is satisfied. If so, the test returns these points and $\tau(A, B) \leq \gamma$ is verified. Otherwise, the test asserts that no pairs satisfy (2.3) and so (2.4) instead must hold. Using this procedure, Gu proposed a bisection-like scheme to estimate $\tau(A, B)$ to *within a factor of two*. For initialization, $\gamma := f(0, 0)$ and $\eta := \gamma$. If the test verifies that $\tau(A, B) \leq \gamma$, then γ and η are both halved (so $\eta = \gamma$ still holds), and the test is done with these smaller values. Otherwise, (2.4) holds, and so γ and $\tau(A, B)$ are within a factor of two and Gu's method terminates.

As noted in [9, p. 358], it is tempting to try to obtain $\tau(A, B)$ to higher precision via a true bisection method, i.e., one that would update both upper and lower bounds, unlike Gu's method, which only updates an upper bound. The problem with this approach is that in order to ascertain whether the current estimate γ is essentially a lower bound to $\tau(A, B)$ via (2.4), one would have to perform the verification procedure for $\eta \approx 0$. Unfortunately, this is not tenable in the presence of rounding errors, as Gu's procedure becomes more and more numerically unreliable as $\eta \rightarrow 0$; i.e., points satisfying (2.3) may not be detected. Consequently, in practice, the lower bound will generally be erroneously updated at some point, thus preventing convergence to $\tau(A, B)$.²

²In Key Remark 6.3, we discuss the unreliability of Gu's procedure in more detail and explain how our new tests avoid key numerical pitfalls. Besides being useful for our $\mathcal{K}(A)$ algorithms, our modifications can also improve the reliability of the $\tau(A, B)$ methods of [14, 9, 15].

Using Gu's verification procedure (as Gu specified it, i.e., without modifications), Burke, Lewis, and Overton instead proposed a trisection algorithm [9, Algorithm 5.2] that balances how much the lower bound is updated with how quickly the value of η is decreased, precisely to postpone the numerical unreliability of Gu's procedure as long as possible. Trisection works as follows. Let $L := 0$ and $U := f(0, 0)$ be initial lower and upper bounds, respectively. Then on the k th iteration, $\eta_k := \frac{2}{3}(U - L)$ and $\gamma_k := L + \eta_k$ are set as the current values of η and γ for the verification test. If the test finds points satisfying (2.3), then the upper bound is updated to $U := \gamma_k$. Otherwise, (2.4) holds, so we know that $\tau(A, B) \geq \gamma_k - \frac{\eta_k}{2} = L + \frac{\eta_k}{2}$, and so now the lower bound can be updated to $L := L + \frac{\eta_k}{2}$. Thus, the new interval has length $\frac{2}{3}(U - L)$. As the trisection algorithm is linearly convergent, Gu's verification procedure will have to be invoked many times, which already is $\mathcal{O}(n^6)$ work by itself using standard dense eigensolvers. Thus, the trisection algorithm also has a large constant factor term hidden away in its asymptotic work complexity. Also, trisection is not a panacea for the numerical issues of Gu's procedure; although $\eta_k \rightarrow 0$ only in the limit, if $\tau(A, B)$ is small, then η_k must become commensurately small in order for trisection to attain any digits of accuracy; see Lemma B.1 and Corollary B.2 in Appendix B.

In the same paper, Burke, Lewis, and Overton also proposed a second algorithm for $\tau(A, B)$ [9, Algorithm 5.3] and advocated it as preferable to trisection. This *optimization-with-restarts* method also relies on Gu's verification procedure, now as a *globality certificate*, and additionally on the fact that $f(x, y)$ is semialgebraic and so $f(x, y)$ only has a finite number of locally minimal values; see [9, p. 359]. This second method thus works by using optimization techniques to find a minimizer of (2.2) with function value f_k and then uses Gu's verification procedure with carefully chosen values of γ and η so that the test checks if f_k is sufficiently close to $\tau(A, B)$; for some relative tolerance $\text{tol} > 0$, the specific values are $\gamma := f_k \cdot (1 - 0.5 \cdot \text{tol})$ and $\eta := f_k \cdot \text{tol}$.³ Otherwise, if $f_k \not\approx \tau(A, B)$ to tolerance, the certificate provides one or more new starting points from which optimization can be restarted with the guarantee that a better (lower) minimum of (2.2) will be found; hence optimization is restarted in a loop until the certificate indeed asserts that the desired accuracy has been attained. By construction, f_k is monotonically decreasing and optimization-with-restarts must terminate with $f_k \approx \tau(A, B)$ to tolerance in a finite number of restarts. Although it is not clear exactly how many restarts will occur, only a handful are typically needed in practice, if any. Furthermore, the optimization phases are relatively cheap, requiring $\mathcal{O}(n^3)$ work with a relatively low constant factor, since minimizers of $f(x, y)$ can generally be found with superlinear or even quadratic convergence. As a result, optimization-with-restarts is almost always many times faster than trisection. However, optimization-with-restarts can still be susceptible to the numerical difficulties of Gu's verification procedure, since η itself may still become very small, e.g., if either high accuracy is desired or $\tau(A, B)$ is small.

Finally, to address the high cost of Gu's verification procedure, though not necessarily its numerical issues, [15] proposed a divide-and-conquer strategy that lowers the asymptotic work complexity of Gu's procedure to $\mathcal{O}(n^4)$ on average and $\mathcal{O}(n^5)$ in the worst case. This benefits all of the aforementioned algorithms.

3. Variational properties and the inverse of the Kreiss constant. We now establish some variational properties of Kreiss constants, which in turn show that locally optimal approximations to $\mathcal{K}(A)$ can be efficiently computed via optimization,

³Note that [9] writes these in an equivalent but different form using $\delta_1 = \gamma$ and $\delta_2 = \gamma - \frac{\eta}{2}$.

even if A is large. We also show how the problem of computing $\mathcal{K}(A)$ shares some similarity with computing $\tau(A, B)$. We begin with the following result.

LEMMA 3.1. *The Kreiss constant \mathcal{K} is not always continuous at A , as it may instantaneously jump to/from ∞ . However, \mathcal{K} is continuous at A if $\alpha(A) < 0$ holds (continuous-time case) or $\rho(A) < 1$ holds (discrete-time case).*

Proof. We begin with the second part. If $\alpha(A) < 0$, then $zI - A$ is invertible for all z such that $\operatorname{Re} z > 0$. By continuity of singular values, $\|(zI - A)^{-1}\|$ in (1.8) is continuous at A , and thus so is $\mathcal{K}(A)$. Via an analogous argument with (1.3), the continuity claim also holds for the discrete-time case.

We prove the first part by example. Let $A(\delta) := \begin{bmatrix} -0.5 & 0 \\ 0 & \delta \end{bmatrix}$ for real scalar $\delta \geq 0$. As $A(\delta)$ is always normal and $\alpha(A(0)) = 0$, $\mathcal{K}(A(0)) = 1$ holds. However, as $\alpha(A(\delta)) > 0$ for any $\delta > 0$, $\mathcal{K}(A(\delta)) = \infty$ for any $\delta > 0$. Using the same example with $\delta \geq 1$ for the discrete-time case, we have that $\mathcal{K}(A(1)) = 1$ and $\mathcal{K}(A(\delta)) = \infty$ for all $\delta > 1$. \square

3.1. The continuous-time case. Identifying \mathbb{C} with \mathbb{R}^2 , consider the inverse of the continuous-time Kreiss constant (1.8), i.e.,

$$(3.1) \quad \mathcal{K}(A)^{-1} = \inf_{x>0, y \in \mathbb{R}} \sigma_{\min} \left(\frac{(x + \mathbf{i}y)I - A}{x} \right) = \inf_{x>0, y \in \mathbb{R}} g(x, y),$$

where

$$(3.2) \quad g(x, y) = \sigma_{\min}(G(x, y)) \quad \text{and} \quad G(x, y) = \frac{(x + \mathbf{i}y)I - A}{x}.$$

Like $f(x, y)$, $g(x, y)$ is semialgebraic, and so in the open right half-plane, $g(x, y)$ must have only a finite number of locally minimal function values. We now derive the gradient and Hessian of $g(x, y)$ which will be useful for finding minimizers via quasi-Newton or Newton methods. Although singular values can vary nonsmoothly with respect to matrix entries, they are nevertheless locally Lipschitz, and so this nonsmoothness is confined to a set of measure zero. We first need the first partial derivatives of $G(x, y)$ for $x \neq 0$:

$$(3.3) \quad \frac{\partial G(x, y)}{\partial x} = \frac{xI - ((x + \mathbf{i}y)I - A)}{x^2} = \frac{A - \mathbf{i}yI}{x^2} \quad \text{and} \quad \frac{\partial G(x, y)}{\partial y} = \frac{\mathbf{i}I}{x}.$$

Let (\hat{x}, \hat{y}) be such that $g(\hat{x}, \hat{y}) \neq 0$ is a simple singular value of $G(x, y)$ with associated left and right singular vectors u and v . Then, by standard perturbation theory for singular values, it follows that

$$(3.4) \quad \nabla g(\hat{x}, \hat{y}) = \operatorname{Re} \begin{bmatrix} u^* \frac{\partial G(x, y)}{\partial x} v \\ u^* \frac{\partial G(x, y)}{\partial y} v \end{bmatrix}.$$

Now since $g(x, y) = \sigma_{\min}(G(x, y))$ is also the n th eigenvalue (in descending order) of the $2n \times 2n$ Hermitian matrix

$$(3.5) \quad \begin{bmatrix} 0 & G(x, y) \\ G(x, y)^* & 0 \end{bmatrix},$$

$\nabla^2 g(\hat{x}, \hat{y})$ can be computed by applying Theorem 1.2 to (3.5). Computationally, the necessary first and second partial derivatives of (3.5) can be obtained via the first partials given in (3.3) and the following second partial derivatives:

$$(3.6) \quad \frac{\partial^2 G(x, y)}{\partial x^2} = \frac{-2(A - \mathbf{i}yI)}{x^3}, \quad \frac{\partial^2 G(x, y)}{\partial y^2} = 0, \quad \text{and} \quad \frac{\partial^2 G(x, y)}{\partial x \partial y} = \frac{-\mathbf{i}I}{x^2}.$$

Although the full eigendecomposition of (3.5) is needed, it can actually be constructed more or less for free given the full SVD of $G(\hat{x}, \hat{y})$; see [4, section 2.2] for details. If (\hat{x}, \hat{y}) is additionally a distinct (up to conjugacy) *global* minimizer of $g(x, y)$, then the gradient and Hessian of $\mathcal{K}(A)^{-1}$ with respect to A are equivalent to the gradient and Hessian of $g(\hat{x}, \hat{y})$ with respect to A . The gradient and Hessian of $\mathcal{K}(A)$ are then simply obtained by applying the chain rule for the inverse.

The cost of obtaining $g(\hat{x}, \hat{y})$, its gradient, and its Hessian is $\mathcal{O}(n^3)$, as they can all be computed given the full SVD of $G(\hat{x}, \hat{y})$. Although (3.1) is technically a constrained optimization problem, $g(x, y) \rightarrow \infty$ as x approaches zero from the right, assuming that $\mathbf{i}y$ is not an eigenvalue of A . Thus, just returning ∞ as the value of $g(x, y)$ whenever $x \leq 0$ suffices for using unconstrained optimization solvers to find feasible local/global minimizers of (3.1). Provided $g(x, y)$ is sufficiently smooth about its stationary points, one can expect local quadratic convergence when using a Newton-based optimization method and superlinear convergence with a quasi-Newton method (forgoing the use of the Hessian). Note that scalable methods for computing smallest singular values, e.g., PROPACK [22], can also be used to compute $g(\hat{x}, \hat{y})$ and its associated pair of left and right singular vectors in order to obtain $\nabla g(\hat{x}, \hat{y})$. Thus, by combining such a sparse solver with a quasi-Newton method, one can efficiently obtain locally optimal approximations to Kreiss constants of large-scale matrices.

Remark 3.2. One could also consider using optimization to find maximizers of (1.9), which has the benefit of working with only one optimization variable instead of two. However, computing $\alpha_\varepsilon(A)$ is substantially more expensive than the minimum singular value of a matrix; the quadratically convergent criss-cross algorithm of [8] to compute $\alpha_\varepsilon(A)$ as well as the faster method of [4] require computing all eigenvalues of $2n \times 2n$ matrices, often several times. Moreover, we have just shown how the explicit Hessian of $g(x, y)$ can easily be computed in order to obtain faster convergence of the iterates produced by optimization methods. Finally, for large-scale A matrices, sparse methods for $\sigma_{\min}(A)$ are generally much faster and more reliable than those for approximating $\alpha_\varepsilon(A)$ [16, 20].

That $g(x, y)$ is not so dissimilar to $f(x, y)$ for $\tau(A, B)$ indicates that it might be possible to adapt Gu's verification procedure to develop globality certificates for $g(x, y)$. Combined with the optimization techniques discussed here, this would enable a globally convergent optimization-with-restarts method for Kreiss constants that terminates within a finite number of restarts.

3.2. The discrete-time case. Again identifying \mathbb{C} with \mathbb{R}^2 , but now using polar coordinates, consider the inverse of the discrete-time Kreiss constant (1.3), i.e.,

$$(3.7) \quad \mathcal{K}(A)^{-1} = \inf_{r>1, \theta \in [0, 2\pi)} \sigma_{\min} \left(\frac{re^{i\theta}I - A}{r - 1} \right) = \inf_{r>1, \theta \in [0, 2\pi)} h(r, \theta),$$

where

$$(3.8) \quad h(r, \theta) = \sigma_{\min}(H(r, \theta)) \quad \text{and} \quad H(r, \theta) = \frac{re^{i\theta}I - A}{r - 1}.$$

Naturally $h(r, \theta)$ has the same key properties as $g(x, y)$, i.e., it too is semialgebraic and locally Lipschitz. Thus, $h(r, \theta)$ has a finite number of locally minimal function values, and we can consider using optimization to find minimizers of $h(r, \theta)$. We will need the analogous gradients and Hessian of $h(r, \theta)$; for brevity, we just provide the

first and second partial derivatives of $H(r, \theta)$ for $r \neq 1$ here, which are, respectively,

$$(3.9) \quad \frac{\partial H(r, \theta)}{\partial r} = \frac{(r-1)e^{i\theta}I - (re^{i\theta}I - A)}{(r-1)^2} = \frac{A - e^{i\theta}I}{(r-1)^2} \quad \text{and} \quad \frac{\partial H(r, \theta)}{\partial \theta} = \frac{ire^{i\theta}I}{r-1},$$

and

$$(3.10) \quad \frac{\partial^2 H(r, \theta)}{\partial r^2} = \frac{-2(A - e^{i\theta}I)}{(r-1)^3}, \quad \frac{\partial^2 H(r, \theta)}{\partial \theta^2} = \frac{-re^{i\theta}I}{r-1}, \quad \text{and} \quad \frac{\partial^2 H(r, \theta)}{\partial r \partial \theta} = \frac{-ie^{i\theta}I}{(r-1)^2}.$$

The costs to compute $h(r, \theta)$ along with its gradient and Hessian also remain as described in subsection 3.1 for $g(x, y)$, and the variational results above similarly allow the gradient and Hessian of $\mathcal{K}(A)^{-1}$ or $\mathcal{K}(A)$ to be computed in the discrete-time case. To ensure that optimization returns a feasible minimizer to (3.7), it suffices to return ∞ for the value of $h(r, \theta)$ whenever $r \leq 1$; this is because for $e^{i\theta}$ not an eigenvalue of A , $\lim_{r \rightarrow 1^+} h(r, \theta) \rightarrow \infty$. Thus, locally optimal approximations to discrete-time Kreiss constants can be computed efficiently for small- or large-scale matrices. To develop a globally convergent algorithm, we will need to develop a discrete-time globality certificate.

4. Continuous-time Kreiss constant analogues of Gu’s theorem and their consequences. Before we present our globally convergent iterations for computing Kreiss constants, we first develop analogues of Theorem 2.1 and Corollary 2.2. For the time being, we consider continuous-time $\mathcal{K}(A)$ and begin by considering vertically oriented pairs of points on the γ -level set of $g(x, y)$. For $\tau(A, B)$, Gu considered pairs of level-set points oriented horizontally, but this choice was rather arbitrary. However, as we will soon see, for Kreiss constants the choice of orientation does have important consequences, both theoretically and for our new algorithms.

THEOREM 4.1. *For $A \in \mathbb{C}^{n \times n}$ with $\alpha(A) < 0$, let $\gamma \in [0, 1)$, $\eta \geq 0$, and (x_*, y_*) be a global minimizer of (3.1). If $\mathcal{K}(A)^{-1} \leq \gamma$ and $\eta \in [0, 2x_*(\gamma - \mathcal{K}(A)^{-1})]$, then there exists a pair $x, y \in \mathbb{R}$ with $x > 0$ such that*

$$(4.1) \quad g(x, y) = g(x, y + \eta) = \gamma.$$

COROLLARY 4.2. *For $A \in \mathbb{C}^{n \times n}$ with $\alpha(A) < 0$, let $\gamma \in [0, 1)$, $\eta \geq 0$, and (x_*, y_*) be a global minimizer of (3.1). If there do not exist any pairs $x, y \in \mathbb{R}$ with $x > 0$ such that (4.1) holds, then*

$$(4.2) \quad \mathcal{K}(A)^{-1} > \gamma - \frac{\eta}{2x_*}.$$

To prove Theorem 4.1, we will use the following topology definition and result.

DEFINITION 4.3. *Let $\mathcal{A} \subset \mathbb{C}$ be a bounded open (path) connected set, and let \mathcal{A}^C be its complement. Furthermore, let $\mathcal{B} = \{z \in \mathbb{C} : z \text{ is in a bounded component of } \mathcal{A}^C\}$. Then $\mathcal{A} \cup \mathcal{B}$ is the simply connected hull of \mathcal{A} , which we denote \mathcal{A}^H .*

LEMMA 4.4. *Let $\mathcal{A}, \mathcal{B} \subset \mathbb{C}$ both be bounded open (path) connected sets. If there exist points $b_{\text{in}}, b_{\text{out}} \in \text{bd } \mathcal{B}^H$ such that $b_{\text{in}} \in \mathcal{A}$ and $b_{\text{out}} \in \mathcal{A}^C$, then $\text{bd } \mathcal{A} \cap \text{bd } \mathcal{B} \neq \emptyset$.*

Proof. Since $\text{bd } \mathcal{B}^H \subset \text{bd } \mathcal{B}$, we can assume that b_{out} is in the interior of \mathcal{A}^C , as otherwise the proof is done. As \mathcal{B}^H is a bounded open simply connected set, $\text{bd } \mathcal{B}^H$ must be connected; see, e.g., [7, p. 345]. By way of contradiction, suppose that $\text{bd } \mathcal{A} \cap \text{bd } \mathcal{B} = \emptyset$. Then $\text{bd } \mathcal{B}^H \subset \text{bd } \mathcal{B} \subset \mathcal{A} \cup \text{int}(\mathcal{A}^C)$. However, since $b_{\text{in}} \in \mathcal{A}$, $b_{\text{out}} \in \text{int}(\mathcal{A}^C)$, and \mathcal{A} and $\text{int}(\mathcal{A}^C)$ are both open and nonempty disjoint sets, it follows by definition of connected that $\text{bd } \mathcal{B}^H$ is in fact disconnected, a contradiction. Thus, $\text{bd } \mathcal{A} \cap \text{bd } \mathcal{B} \neq \emptyset$ holds. \square

Proof of Theorem 4.1. If $\gamma = \mathcal{K}(A)^{-1}$, the proof is trivially satisfied with $\eta = 0$, so assume that $\gamma \in (\mathcal{K}(A)^{-1}, 1)$. Since $\alpha(A) < 0$, it follows that $\lim_{x \rightarrow 0^+} g(x, y) = \infty$ for any $y \in \mathbb{R}$, and so $g(x_*, y_*) = \mathcal{K}(A)^{-1}$ with $x_* > 0$. Now consider the strict lower level set $\mathcal{L}_\gamma := \{(x, y) : g(x, y) < \gamma, x > 0\}$, which is clearly open and also bounded; see [27, Theorem 2.3]. Let \mathcal{L} be the (open) connected component of \mathcal{L}_γ such that $(x_*, y_*) \in \mathcal{L}$, and let $\mathcal{G} := \mathcal{L}^H$, i.e., the simply connected hull of \mathcal{L} .

Now, following the proof of [14, Theorem 3.1] a bit more closely, by continuity of $g(x, y)$ and the fact that \mathcal{G} is bounded, there must exist points $b_1, b_2 \in \text{bd } \mathcal{G}$,

$$b_1 = (x_*, y_* - \eta_1) \quad \text{and} \quad b_2 = (x_*, y_* + \eta_2),$$

such that

$$(4.3) \quad g(x_*, y_* - \eta_1) = g(x_*, y_* + \eta_2) = \gamma,$$

where $\eta_1, \eta_2 > 0$. Furthermore, we can assume that η_1 and η_2 are the smallest positive values such that (4.3) holds with $b_1, b_2 \in \text{bd } \mathcal{G}$. Noting that

$$(4.4) \quad g(x, y) = \frac{\sigma_{\min}((x + \mathbf{i}y)I - A)}{x} = \frac{f(x, y)}{x},$$

whose numerator has a GLC of 1, it follows for any $y_1, y_2 \in \mathbb{R}$ that

$$|g(x_*, y_1) - g(x_*, y_2)| \leq \frac{1}{x_*} |y_1 - y_2|;$$

i.e., $g(x_*, y)$ with respect to y has a GLC of $\frac{1}{x_*}$. By applying this GLC to (4.3), it follows that

$$(4.5) \quad \eta_1 \geq x_*(\gamma - \mathcal{K}(A)^{-1}) \quad \text{and} \quad \eta_2 \geq x_*(\gamma - \mathcal{K}(A)^{-1}),$$

and so $\eta_1 + \eta_2 \geq 2x_*(\gamma - \mathcal{K}(A)^{-1})$.

Now suppose that $\eta \in (0, 2x_*(\gamma - \mathcal{K}(A)^{-1})]$ so $\eta \leq \eta_1 + \eta_2$. Obviously (4.1) is satisfied if $\eta = \eta_1 + \eta_2$, so assume that $\eta < \eta_1 + \eta_2$. Let

$$\mathcal{G}_\eta := \{(x, y - \eta) : (x, y) \in \mathcal{G}\},$$

i.e., \mathcal{G} shifted downward by the amount η , and consider the line segment joining b_1 and b_2 and the point $b_{\text{in}} = (x_*, y_* + \eta_2 - \eta)$. As b_{in} must be on this line segment, but not at its endpoints, $b_{\text{in}} \in \mathcal{G}$ and $b_{\text{in}} \in \text{bd } \mathcal{G}_\eta$, since $b_2 \in \text{bd } \mathcal{G}$. Let $b_{\text{out}} \in \text{bd } \mathcal{G}_\eta$ be a lowermost point of $\text{bd } \mathcal{G}_\eta$. Then, as $\eta > 0$, $b_{\text{out}} \in \mathcal{G}^C$, the complement of G . Since \mathcal{G} and \mathcal{G}_η are both bounded open connected sets in the plane and $\mathcal{G}_\eta = \mathcal{G}_\eta^H$, Lemma 4.4 applies, and so $\text{bd } \mathcal{G} \cap \text{bd } \mathcal{G}_\eta \neq \emptyset$. Letting (\tilde{x}, \tilde{y}) be any such point in $\text{bd } \mathcal{G} \cap \text{bd } \mathcal{G}_\eta$, it follows that $(\tilde{x}, \tilde{y} + \eta) \in \text{bd } \mathcal{G}$; hence (\tilde{x}, \tilde{y}) satisfies (4.1). \square

We now consider horizontally oriented pairs of points on the γ -level set of $g(x, y)$, similar to Theorem 2.1 for $f(x, y)$. As the horizontal orientation is actually more complicated for $\mathcal{K}(A)$, we first establish the following intermediate general result.

LEMMA 4.5. *Let $f : (0, \infty) \rightarrow [0, \infty)$ be continuous with a GLC of $c \geq 0$, and consider the function $q(x) := \frac{f(x)}{x}$ on the same domain. For $a, b > 0$ with $b - a = \eta > 0$, if $q(a) = q(b) = \gamma$ and $x_* = \arg \min_{x \in [a, b]} q(x)$ with $\gamma_* = q(x_*)$, then*

$$\gamma_* \geq \gamma - \frac{\eta(c + \gamma)}{2x_*}.$$

Proof. We assume that $\gamma_* < \gamma$, as otherwise the inequality clearly holds. Since $f(x)$ has a GLC of c , it follows that

$$f(a) - f(x_*) + f(b) - f(x_*) \leq c(x_* - a) + c(b - x_*) = c(b - a) = c\eta.$$

Meanwhile

$$f(a) - f(x_*) + f(b) - f(x_*) = a\gamma + b\gamma - 2x_*\gamma_* \geq (a + 2x_* - b)\gamma - 2x_*\gamma_* = 2x_*(\gamma - \gamma_*) - \eta\gamma.$$

Combining the two yields $2x_*(\gamma - \gamma_*) - \eta\gamma \leq c\eta$, thus completing the proof. \square

THEOREM 4.6. For $A \in \mathbb{C}^{n \times n}$ with $\alpha(A) < 0$, let $\gamma \in [0, 1)$, $\eta \geq 0$, and (x_*, y_*) be a global minimizer of (3.1). If $\mathcal{K}(A)^{-1} \leq \gamma$ and $\eta \in [0, \frac{2x_*}{1+\gamma}(\gamma - \mathcal{K}(A)^{-1})]$, then there exists a pair $x, y \in \mathbb{R}$ with $x > 0$ such that

$$(4.6) \quad g(x, y) = g(x + \eta, y) = \gamma.$$

COROLLARY 4.7. For $A \in \mathbb{C}^{n \times n}$ with $\alpha(A) < 0$, let $\gamma \in [0, 1)$, $\eta \geq 0$, and (x_*, y_*) be a global minimizer of (3.1). If there do not exist any pairs $x, y \in \mathbb{R}$ with $x > 0$ such that (4.6) holds, then

$$(4.7) \quad \mathcal{K}(A)^{-1} > \gamma - \frac{\eta(1+\gamma)}{2x_*}.$$

Proof of Theorem 4.6. The beginning of the proof is the same as the first paragraph of the proof of Theorem 4.1. Again consider the simply connected set \mathcal{G} defined there that contains (x_*, y_*) , a global minimizer. By continuity of $g(x, y)$ and the fact that \mathcal{G} is bounded, there must exist points $b_1, b_2 \in \text{bd } \mathcal{G}$,

$$b_1 = (x_* - \eta_1, y_*) \quad \text{and} \quad b_2 = (x_* + \eta_2, y_*),$$

such that

$$(4.8) \quad g(x_* - \eta_1, y_*) = g(x_* + \eta_2, y_*) = \gamma,$$

where $\eta_1, \eta_2 > 0$. We again assume that η_1 and η_2 are the smallest positive values such that (4.8) holds with $b_1, b_2 \in \text{bd } \mathcal{G}$. Applying Lemma 4.5 to $q(x) := g(x, y_*) = \frac{f(x, y_*)}{x}$ with $a = x_* - \eta_1$ and $b = x_* + \eta_2$, we have that

$$\eta_1 + \eta_2 \geq \frac{2x_*}{1+\gamma}(\gamma - \mathcal{K}(A)^{-1}),$$

as the numerator of $g(x, y)$ as rewritten in (4.4) has a GLC of 1.

Now suppose that $\eta \in (0, \frac{2x_*}{1+\gamma}(\gamma - \mathcal{K}(A)^{-1})]$ so $\eta \leq \eta_1 + \eta_2$. If $\eta = \eta_1 + \eta_2$, (4.6) is clearly satisfied, so instead assume that $\eta < \eta_1 + \eta_2$. Considering the set $\widehat{\mathcal{G}} := \{(x - \eta, y) : (x, y) \in \mathcal{G}\}$, i.e., \mathcal{G} shifted to the left by the amount η , the rest of the proof follows analogously to the end of the proof of Theorem 4.1. \square

Although we have derived Kreiss constant analogues of Theorem 2.1 and Corollary 2.2, Corollaries 4.2 and 4.7 in fact assert that the $\tau(A, B)$ algorithms of [14, 9, 15] will *not* directly extend to Kreiss constants. The crux of the problem is that these $\tau(A, B)$ methods all rely on the fact that (2.4) holds when there are no points satisfying (2.3), which, recall, provided a way of computing $\tau(A, B)$ via Gu's verification procedure to verify an upper or lower bound for $\tau(A, B)$. However, in the Kreiss constant setting, our lower bounds given in (4.2) and (4.7) are not as concrete, as they depend on x_* , which is unknown. For a given γ and η , we do not even know if the lower bounds provided in (4.2) and (4.7) would be meaningful, as they might not even be positively valued. Thus, to develop algorithms for $\mathcal{K}(A)$, crucial departures must be made. We will do this via two different strategies.

Algorithm 5.1. Optimization-with-restarts using backtracking.

Input: A nonnormal matrix $A \in \mathbb{C}^{n \times n}$ with $\alpha(A) < 0$, $x_0 > 0$, and $y_0 \in \mathbb{R}$ such that $g(x_0, y_0) < 1$, and a tolerance $\eta_{\text{tol}} > 0$.

Output: $\gamma^{-1} \approx \mathcal{K}(A)$ (continuous-time).

```

1: while true do
2:    $(\hat{x}, \hat{y}) \leftarrow$  computed local/global minimizer of (3.1) initialized from  $(x_0, y_0)$ 
3:    $\gamma \leftarrow g(\hat{x}, \hat{y})$ 
4:    $\eta \leftarrow$  some positive value  $\gg \eta_{\text{tol}}$ 
5:   while true do
6:     Perform the 2D level-set test of subsection 5.2 with current  $\gamma$  and  $\eta$ 
7:     if test finds any level-set points then
8:        $(x_0, y_0) \leftarrow$  one of these points
9:       break // Goto line 2 to restart optimization.
10:    else if  $\eta \leq \eta_{\text{tol}}$  then
11:      return // Found a global minimizer to tolerance.
12:    else
13:       $\eta \leftarrow c\eta$  for some constant  $c \in (0, 1)$ 
14:    end if
15:  end while
16: end while

```

NOTE: For simplicity of the pseudocodes, we assume here and in Algorithm 6.1 that (a) optimization always converges to local or global minimizers exactly, i.e., not approximately or to other stationary points, and (b) points found by the certificate test (if any) are never exactly stationary.

5. A continuous-time $\mathcal{K}(A)$ algorithm based on fixed-distance pairs.

Let (\hat{x}, \hat{y}) be a local (but not global) minimizer of $g(x, y)$, and set $\gamma := g(\hat{x}, \hat{y})$. As (\hat{x}, \hat{y}) is a local minimizer, we do not need to verify that $\gamma \geq \mathcal{K}(A)^{-1}$, as this is obviously true. To obtain an optimization-with-restarts algorithm, we do not necessarily need to verify a lower bound either. Instead, we can just aim to detect other (non-stationary) points on the γ -level set of $g(x, y)$. Using such level-set points to restart optimization, a better (lower) minimizer of $g(x, y)$ is guaranteed to be found. Assuming $\gamma \in [\mathcal{K}(A)^{-1}, 1)$, Theorems 4.1 and 4.6 assert that if $\eta > 0$ is chosen sufficiently small, there must exist points satisfying either (4.1) or (4.6). Of course, we do not know a priori how small to choose η , so we propose using backtracking; i.e., we can start with η initially set to some large value and simply decrease it in a loop until level-set points for restarting optimization are found. As long as (\hat{x}, \hat{y}) is not a global minimizer, this backtracking procedure must succeed in finding level-set points for restarting optimization. Meanwhile, when (\hat{x}, \hat{y}) is a global minimizer, and so $\mathcal{K}(A)$ has been computed, the backtracking procedure can simply be terminated once η falls below a tolerance. A high-level pseudocode using this backtracking-based globality certificate is given in Algorithm 5.1. To complete this algorithm, we now must develop a corresponding 2D level-set test for continuous-time Kreiss constants. We will do this by looking for points satisfying (4.1) or (4.6), i.e., level-set points that are a fixed distance η apart, and develop a procedure inspired by Gu's 2D level-set test for $\tau(A, B)$.

5.1. A 1D vertical level-set test. Before we develop our 2D level-set test for $g(x, y)$, we will need the following theorem, which will allow us to obtain all the points on the γ -level set of $g(x, y)$ along a chosen vertical line.

THEOREM 5.1. *Given $\gamma, x, y \in \mathbb{R}$, with $\gamma \geq 0$ and $x \neq 0$, γ is a singular value of $G(x, y)$ defined in (3.2) if and only if $\mathbf{i}y$ is an eigenvalue of the Hamiltonian matrix*

$$(5.1) \quad \begin{bmatrix} A - xI & \gamma xI \\ -\gamma xI & xI - A^* \end{bmatrix}.$$

Proof. It is clear that the matrix is Hamiltonian. Suppose γ is a singular value $G(x, y)$ with left and right singular vectors u and v , and so

$$\gamma \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} G(x, y) & 0 \\ 0 & G(x, y)^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix} \Leftrightarrow \gamma x \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} (x + \mathbf{i}y)I - A & 0 \\ 0 & (x - \mathbf{i}y)I - A^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix}.$$

Rearranging terms and multiplying the bottom block row by -1 , this is equivalent to

$$\gamma x \begin{bmatrix} u \\ -v \end{bmatrix} + \begin{bmatrix} A - xI & 0 \\ 0 & xI - A^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix} = \mathbf{i}y \begin{bmatrix} v \\ u \end{bmatrix} \Leftrightarrow \begin{bmatrix} A - xI & \gamma xI \\ -\gamma xI & xI - A^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix} = \mathbf{i}y \begin{bmatrix} v \\ u \end{bmatrix}. \quad \square$$

In fact, for a given $\gamma \geq \mathcal{K}(A)^{-1}$ and $x := \hat{x} \neq 0$, computing the imaginary eigenvalues of (5.1) may provide more than the γ -level set points of $g(x, y)$ along the vertical line $x = \hat{x}$. This is because if $\mathbf{i}\hat{y}$ is an eigenvalue of (5.1), Theorem 5.1 asserts that γ is a singular value of $F(\hat{x}, \hat{y})$, but not necessarily the minimum one; if this happens, (\hat{x}, \hat{y}) would be on a level set of $g(x, y)$ lower than the γ -level set and thus be an even better point for restarting optimization. Finally, note that structure-preserving eigensolvers [2, 3] can be used for more reliable detection of imaginary eigenvalues of Hamiltonian matrices like (5.1).

5.2. A 2D level-set test for fixed-distance pairs. We now derive a new 2D level-set test for $g(x, y)$. Per Theorems 4.1 and 4.6, the choice of orientation for pairs of points on the γ -level set of $g(x, y)$ has consequences; hence we will develop our new continuous-time Kreiss constant procedure for arbitrary orientation. Specifically, given $\eta > 0$ and angle $\theta \in (-\frac{\pi}{2}, \frac{\pi}{2}]$, we will look for points a fixed distance η apart of the form (\hat{x}, \hat{y}) and $(\hat{x} + \eta \cos \theta, \hat{y} + \eta \sin \theta)$ such that $g(x, y) = \gamma$ holds at both of them and $\hat{x} > 0$.

Suppose that γ is a singular value of both $G(x, y)$ and $G(x + \eta \cos \theta, y + \eta \sin \theta)$, with respective left and right singular vectors pairs u, v and \hat{u}, \hat{v} . By Theorem 5.1 for $G(x, y)$ and following a similar argument as in its proof for $G(x + \eta \cos \theta, y + \eta \sin \theta)$, the following two Hamiltonian matrices must share an imaginary eigenvalue $\mathbf{i}y$:

$$(5.2a) \quad \begin{bmatrix} A - xI & \gamma xI \\ -\gamma xI & xI - A^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix} = \mathbf{i}y \begin{bmatrix} v \\ u \end{bmatrix},$$

$$(5.2b) \quad \begin{bmatrix} A - (x + \eta e^{\mathbf{i}\theta})I & \gamma(x + \eta \cos \theta)I \\ -\gamma(x + \eta \cos \theta)I & (x + \eta e^{-\mathbf{i}\theta})I - A^* \end{bmatrix} \begin{bmatrix} \hat{v} \\ \hat{u} \end{bmatrix} = \mathbf{i}y \begin{bmatrix} \hat{v} \\ \hat{u} \end{bmatrix}.$$

Let A_1 and A_2 respectively denote the square matrices in (5.2a) and (5.2b), and let $W = \begin{bmatrix} v \\ u \end{bmatrix} \begin{bmatrix} \hat{v}^* & \hat{u}^* \end{bmatrix} \neq 0$ so that we have $A_1 W = \mathbf{i}y W$ and $A_2 W^* = \mathbf{i}y W^*$. To eliminate y , we take the conjugate transpose of the second and then add the two together to obtain the Sylvester equation:

$$(5.3) \quad \begin{bmatrix} A - xI & \gamma xI \\ -\gamma xI & xI - A^* \end{bmatrix} W + W \begin{bmatrix} A^* - (x + \eta e^{-\mathbf{i}\theta})I & -\gamma(x + \eta \cos \theta)I \\ \gamma(x + \eta \cos \theta)I & (x + \eta e^{\mathbf{i}\theta})I - A \end{bmatrix} = 0.$$

Thus, if (5.3) also has a nonzero solution $W \in \mathbb{C}^{2n \times 2n}$, A_1 and $-A_2^*$ must share an eigenvalue. As A_1 and A_2 are Hamiltonian matrices, their spectra have imaginary-axis symmetry, and so A_1 and A_2 must also have an eigenvalue in common. Now, separating out all terms involving x , we get

$$(5.4) \quad \left(\begin{bmatrix} A & 0 \\ 0 & -A^* \end{bmatrix} W + W \begin{bmatrix} A^* - \eta e^{-i\theta} I & -\gamma \eta \cos \theta I \\ \gamma \eta \cos \theta I & \eta e^{i\theta} I - A \end{bmatrix} \right) - x \left(\begin{bmatrix} I & -\gamma I \\ \gamma I & -I \end{bmatrix} W + W \begin{bmatrix} I & \gamma I \\ -\gamma I & -I \end{bmatrix} \right) = 0.$$

Rewriting both Sylvester forms using the vectorize operator, and letting $w = \text{vec}(W)$, results in the generalized eigenvalue problem

$$(5.5) \quad \mathcal{A}_1 w = x \mathcal{A}_2 w, \quad \text{where}$$

$$\mathcal{A}_1 = I_{2n} \otimes \begin{bmatrix} A & 0 \\ 0 & -A^* \end{bmatrix} + \begin{bmatrix} \bar{A} - \eta e^{-i\theta} I & \gamma \eta \cos \theta I \\ -\gamma \eta \cos \theta I & \eta e^{i\theta} I - A^\top \end{bmatrix} \otimes I_{2n},$$

$$\mathcal{A}_2 = I_{2n} \otimes \begin{bmatrix} I & -\gamma I \\ \gamma I & -I \end{bmatrix} + \begin{bmatrix} I & -\gamma I \\ \gamma I & -I \end{bmatrix} \otimes I_{2n}.$$

Our 2D level-set test for fixed-distance pairs works as follows. Given estimate $\gamma \in [\mathcal{K}(A)^{-1}, 1)$, $\eta \geq 0$, and $\theta \in (-\frac{\pi}{2}, \frac{\pi}{2}]$, we first compute all the eigenvalues of (5.5). If there are no (finite) positive real eigenvalues of (5.5), then the test is finished and returns no level-set points. Otherwise, additional calculations are done to ascertain whether or not any level-set points have been detected. For each eigenvalue $\hat{x} > 0$ of (5.5), we have that the vertical line specified by \hat{x} , and, if $|\theta| < \frac{\pi}{2}$, the vertical line $\hat{x} + \eta \cos \theta$, may contain points on the γ -level set of $g(x, y)$. To determine this, for each of these vertical lines, say, $x := \hat{x}$, we then apply Theorem 5.1 and compute all the eigenvalues of the corresponding Hamiltonian matrix (5.1). If this matrix has no imaginary eigenvalues, then no level-set points have been detected on this vertical line. Otherwise, all the points (\hat{x}, \hat{y}) such that $i\hat{y}$ is an imaginary eigenvalue of this matrix are added to the list of detected level-set points of $g(x, y)$ to return. Optionally, for each of these points, one could additionally check whether or not γ is the minimum singular value of $G(\hat{x}, \hat{y})$, but this is not strictly necessary; as discussed previously, (\hat{x}, \hat{y}) must be on the γ -level set of $g(x, y)$ or a lower one, either of which suffices for restarting optimization to obtain a better (lower) minimizer of $g(x, y)$. If no level-set points are detected, for any of the vertical lines, then the test returns no points. Otherwise, all the detected level-set points are returned.

Our new 2D level-set test differs from the procedure in [14] (and [15]) in a significant way; in Key Remark 6.3, we explain how our modifications here greatly improve the reliability of these 2D level-set tests.

In terms of cost, in extreme situations there may be up to $\mathcal{O}(n^2)$ potential vertical lines detected, which means that $\mathcal{O}(n^2)$ Hamiltonian eigenvalue problems of dimension $2n \times 2n$ must be solved. However, when using standard dense eigensolvers, the overall work complexity of our procedure is actually $\mathcal{O}(n^6)$, as we must first compute the eigenvalues of (5.5), which is a matrix pencil with square matrices of dimension $4n^2$. In terms of constant factors, if A is real, \mathcal{A}_1 is real if and only if $\theta = 0$.

5.3. Properties of the eigensystem $\mathcal{A}_1 w = x \mathcal{A}_2 w$ and its solution. Unlike the $4n^2 \times 4n^2$ generalized eigenvalue problem (5.5), Gu derived a smaller $2n^2 \times 2n^2$ generalized eigenvalue problem for $\tau(A, B)$ [14, eq. (3.13)]. In [15, section 3.1], this

was then simplified further to a computationally easier $2n^2 \times 2n^2$ standard eigenvalue problem [15, eq. (3.7)]. Via the following result, we show that the rank of \mathcal{A}_2 is $2n^2$, i.e., half its dimension.

LEMMA 5.2. *Let $C := \begin{bmatrix} aI & -bI \\ bI & -aI \end{bmatrix} \in \mathbb{C}^{2n \times 2n}$ with $a, b \in \mathbb{C}$ and $b \neq 0$. For $k \in \mathbb{N}$, the Kronecker sum $I_{2k} \otimes C + C \otimes I_{2k} = \mathcal{U}_k \mathcal{V}_k \in \mathbb{C}^{4kn \times 4kn}$ and $\mathcal{V}_k \mathcal{U}_k = 2I_k \otimes C$, where*

$$(5.6a) \quad \mathcal{U}_k := \begin{bmatrix} I_k \otimes \begin{bmatrix} 2aI & -bI \\ bI & 0 \end{bmatrix} & \\ & bI_{2kn} \end{bmatrix} \in \mathbb{C}^{4kn \times 2kn},$$

$$(5.6b) \quad \mathcal{V}_k := \begin{bmatrix} I_{2kn} & | & I_k \otimes \begin{bmatrix} 0 & -I \\ I & -2ab^{-1}I \end{bmatrix} \end{bmatrix} \in \mathbb{C}^{2kn \times 4kn}.$$

Proof. The factorization follows from the following if-and-only-if equivalences:

$$(5.7) \quad \begin{aligned} \mathcal{U}_k \mathcal{V}_k &= \begin{bmatrix} I_k \otimes \begin{bmatrix} 2aI & -bI \\ bI & 0 \end{bmatrix} & | & -bI_{2kn} \\ & & & I_k \otimes \begin{bmatrix} 0 & -bI \\ bI & -2aI \end{bmatrix} \end{bmatrix} \\ &= \begin{bmatrix} I_k \otimes \begin{bmatrix} aI & -bI \\ bI & -aI \end{bmatrix} & | & 0 \\ & & & I_k \otimes \begin{bmatrix} aI & -bI \\ bI & -aI \end{bmatrix} \end{bmatrix} + \begin{bmatrix} aI_{2kn} & | & -bI_{2kn} \\ bI_{2kn} & | & -aI_{2kn} \end{bmatrix} \\ &= I_{2k} \otimes \begin{bmatrix} aI & -bI \\ bI & -aI \end{bmatrix} + \begin{bmatrix} aI & -bI \\ bI & -aI \end{bmatrix} \otimes I_{2k}, \end{aligned}$$

while $\mathcal{V}_k \mathcal{U}_k = I_k \otimes \begin{bmatrix} 2aI & -bI \\ bI & 0 \end{bmatrix} + bI_k \otimes \begin{bmatrix} 0 & -I \\ I & -2ab^{-1}I \end{bmatrix} = 2I_k \otimes C. \quad \square$

Applying Lemma 5.2 to \mathcal{A}_2 with $a := 1, b := \gamma \neq 0$, and $k := n, \mathcal{V}_k$ gives the reduced row echelon form of \mathcal{A}_2 , and so the rank of \mathcal{A}_2 is $2n^2$. Thus, is it natural to ask if (5.5) can also be reduced to a $2n^2 \times 2n^2$ generalized eigenvalue problem, and perhaps even a standard one. Unfortunately, the presence of nonzero off-diagonal blocks $\pm\gamma I$ in $\begin{bmatrix} I & -\gamma I \\ \gamma I & -I \end{bmatrix}$ from \mathcal{A}_2 appears to prevent this. If one attempts to follow [14, 15] and similarly partition W into four $n \times n$ blocks, multiplying out (5.4) for each block of W results in four equations that all involve x . In contrast, in [14, 15], eigenvalue x (α in their notation) only appears in the two corresponding equations for the diagonal blocks of W (X in their notation); for the off-diagonal blocks, eigenvalue x (again α in their notation) does not appear in these other two equations, since, as seen in [15, eq. (3.5)], it ends up being multiplied by zero. Consequently, the reduction techniques of [14, 15] do not seem to be applicable to (5.5).

However, since \mathcal{A}_2 is singular, $\mathcal{A}_1 - \lambda\mathcal{A}_2$ can at least be *numerically* deflated into a smaller pencil $\tilde{\mathcal{A}}_1 - \lambda\tilde{\mathcal{A}}_2$ whose spectrum is the set of finite eigenvalues of (5.5); see the deflation routines of [6, section 4.3] and [18, Algorithm 3, Chapter 3.1], the former of which is implemented in the `ml_ct_dss_adtf` routine from MORLAB [5]. Although this iterative deflation technique is cubic work, and so also $\mathcal{O}(n^6)$ work for (5.5), deflating matrix pencils this way can be faster than computing their eigenvalues with the QZ algorithm. Furthermore, for (5.5), deflation results in a matrix pencil of half the order, since it removes $2n^2$ infinite eigenvalues. As a result, deflating and then computing the eigenvalues of the resulting pencil may be even faster than computing the eigenvalues of (5.5) directly. That matrix $\tilde{\mathcal{A}}_2$ in the reduced pencil is nonsingular also provides a second important benefit. Although we must compute the real eigenvalues of (5.5), we cannot expect its real eigenvalues to be exactly real in the presence of rounding errors. The key question then is how far away from the real axis can a computed eigenvalue be allowed to be while still being *deemed* a real eigenvalue of (5.5)? For the reduced pencil, a reliable tolerance is `tol` $\cdot \epsilon_{\text{mach}} \cdot \|\tilde{\mathcal{A}}_2^{-1} \tilde{\mathcal{A}}_1\|_\infty$, where

$\text{tol} > 1$ is provided by the user and ϵ_{mach} is the machine precision. Note that we still recommend computing the eigenvalues of $\tilde{\mathcal{A}}_1 - \lambda\tilde{\mathcal{A}}_2$ as a generalized eigenvalue problem, instead of using the matrix $\tilde{\mathcal{A}}_2^{-1}\tilde{\mathcal{A}}_1$, as we have observed that the condition number of $\tilde{\mathcal{A}}_2$ is generally very large in practice.

5.4. Faster computation of the real eigenvalues of $\mathcal{A}_1 w = x\mathcal{A}_2 w$. We now show how the $\mathcal{O}(n^6)$ theoretical work complexity of Algorithm 5.1 can be reduced. To do this, we will work with the matrices in (5.5) and adapt the divide-and-conquer approach proposed in [15, section 3.3.2] for faster computation of $\tau(A, B)$. At a high level, this efficiency improvement relies on two principles. First, although the matrices in (5.5) are $4n^2 \times 4n^2$ in size, they arose from vectorizing the two corresponding $2n \times 2n$ Sylvester forms in (5.4). As such, applying \mathcal{A}_1 and \mathcal{A}_2 or their inverses to a vector can actually be done with just $\mathcal{O}(n^3)$ work. In turn, this means that for any shift $s \in \mathbb{C}$, $(\mathcal{A}_1 - s\mathcal{A}_2)^{-1}$ can be applied to a vector with $\mathcal{O}(n^3)$ work (we will clarify how these computations are done in a moment). Consequently, a shift-and-invert eigenvalue solver, e.g., `eigs` in MATLAB, can be employed to find the eigenvalues of (5.5) that are closest to a shift s with only $\mathcal{O}(n^3)$ work. Second, given a matrix $\mathcal{X} \in \mathbb{C}^{q \times q}$, suppose one only wants its eigenvalues that are along a line segment, say, an interval $[0, D]$ on the real axis for some $D > 0$. Then the recursive iteration given by [15, Algorithm 4], which uses a shift-and-invert eigensolver, can locate all eigenvalues of \mathcal{X} in $[0, D]$ with at most $2q + 1$ shifts in the worst case and $\mathcal{O}(\sqrt{q})$ shifts if the eigenvalues of \mathcal{X} are distributed uniformly. For brevity, we forgo the details of describing [15, Algorithm 4] but note that the cost of choosing D unnecessarily large only results in about four extra shifts [15, p. 490]. Thus, the theory says that by adapting this divide-and-conquer technique to compute the positive real eigenvalues of (5.5), the overall work complexity of Algorithm 5.1 will be reduced to $\mathcal{O}(n^4)$ work on average and $\mathcal{O}(n^5)$ in the worst case.

We now explain how \mathcal{A}_1 , \mathcal{A}_2 , and $(\mathcal{A}_1 - s\mathcal{A}_2)^{-1}$ can all be applied to a vector $w \in \mathbb{C}^{4n^2}$ with at most $\mathcal{O}(n^3)$ work. As \mathcal{A}_2 has only $10n^2$ nonzero entries, it suffices to store it in a sparse matrix format. Computing $\mathcal{A}_1 w$ can be done efficiently via vectorizing the first Sylvester form in (5.4), i.e., $\text{vec} \left(\begin{bmatrix} A & 0 \\ 0 & -A^* \end{bmatrix} W + W \begin{bmatrix} A^* - \eta e^{-i\theta} I - \gamma \eta \cos \theta I & \\ \gamma \eta \cos \theta I & \eta e^{i\theta} I - A \end{bmatrix} \right)$, where $W \in \mathbb{C}^{2n \times 2n}$ and $w = \text{vec}(W)$. The dominant cost in obtaining $\mathcal{A}_1 w$ is the two matrix multiplications with W ; hence it too can be done in $\mathcal{O}(n^3)$ work. For $y \in \mathbb{C}^{4n^2}$, we can efficiently obtain $w = (\mathcal{A}_1 - s\mathcal{A}_2)^{-1}y$ by considering $(\mathcal{A}_1 - s\mathcal{A}_2)w = y$. This “unvectorizes” into (5.3), provided that x is replaced with s and the zero in its right-hand side is replaced by Y , where $y = \text{vec}(Y)$. By solving the resulting Sylvester equation and vectorizing its solution W , $w = (\mathcal{A}_1 - s\mathcal{A}_2)^{-1}y$ is computed in $\mathcal{O}(n^3)$ work.

6. Continuous-time $\mathcal{K}(A)$ algorithms based on variable-distance pairs.

Having developed the first globally convergent iteration for continuous-time Kreiss constants, we now develop two more, namely Algorithms 6.1 and 6.2, which can be considered closer analogues of the two $\tau(A, B)$ methods of [9] described in section 2. As previously discussed at the end of section 4, the lower bounds provided by Corollaries 4.2 and 4.7, due to their dependence on the unknown x_* , prevent creating direct $\mathcal{K}(A)$ analogues of the trisection and optimization-with-restarts of [9]. However, these lower bounds are the result of the assumption of looking for pairs of level-set points that are a *fixed distance* η apart, as Gu originally used for Theorem 2.1. As we are about to show, if we change this assumption, then we can obtain *different* lower bound results than those given in (4.2) and (4.7).

Algorithm 6.1. Optimization-with-restarts (no backtracking).

Input: A nonnormal matrix $A \in \mathbb{C}^{n \times n}$ with $\alpha(A) < 0$, $x_0 > 0$, and $y_0 \in \mathbb{R}$ such that $g(x_0, y_0) < 1$, and a tolerance $\gamma_{\text{tol}} > 0$.

Output: $g_k^{-1} \approx \mathcal{K}(A)$ (continuous-time).

```

1: while true do
2:    $(\hat{x}, \hat{y}) \leftarrow$  computed local/global minimizer of (3.1) initialized from  $(x_0, y_0)$ 
3:    $g_k \leftarrow g(\hat{x}, \hat{y})$ 
4:    $\gamma \leftarrow g_k(1 - 0.5 \cdot \gamma_{\text{tol}})$ 
5:    $\eta \leftarrow g_k \cdot \gamma_{\text{tol}}$ 
6:   Perform the 2D level-set test of subsection 6.1 with current  $\gamma$  and  $\eta$ 
7:   if test finds any level-set points then
8:      $(x_0, y_0) \leftarrow$  one of these points // Goto line 2 to restart optimization.
9:   else
10:    return //  $\mathcal{K}(A)^{-1} > g_k \cdot (1 - \gamma_{\text{tol}})$  holds.
11:  end if
12: end while

```

THEOREM 6.1. For $A \in \mathbb{C}^{n \times n}$ with $\alpha(A) < 0$, let $\gamma \in [0, 1)$, $\eta \geq 0$, and (x_*, y_*) be a global minimizer of (3.1). If $\mathcal{K}(A)^{-1} \leq \gamma$ and $\eta \in [0, 2(\gamma - \mathcal{K}(A)^{-1})]$, then there exists a pair $x, y \in \mathbb{R}$ with $x > 0$ such that

$$(6.1) \quad g(x, y) = g(x, y + x\eta) = \gamma.$$

COROLLARY 6.2. For $A \in \mathbb{C}^{n \times n}$ with $\alpha(A) < 0$, let $\gamma \in [0, 1)$, $\eta \geq 0$, and (x_*, y_*) be a global minimizer of (3.1). If there do not exist any pairs $x, y \in \mathbb{R}$ with $x > 0$ such that (6.1) holds, then

$$(6.2) \quad \mathcal{K}(A)^{-1} > \gamma - \frac{\eta}{2x_*}.$$

Proof of Theorem 6.1. The proof follows similarly to the proof of Theorem 4.1 with the following modification. By using the variable distance $x\eta$, for $\gamma \in [\mathcal{K}(A)^{-1}, 1)$ we instead obtain that $\eta \in [0, 2(\gamma - \mathcal{K}(A)^{-1})]$ is a sufficient condition for (6.1) to hold; this is because this choice cancels out the x_* in the proof, as $|y_1 - y_2| = x_*\eta$. \square

Thus, by looking for vertical pairs of points that are this particular *variable distance* apart, i.e., $x\eta$, a corresponding certificate procedure would assert either that $\gamma \geq \mathcal{K}(A)^{-1}$ holds or that (6.2) does. Such a certificate would avoid the need for the backtracking procedure that was necessary for Algorithm 5.1. While this might seem to be obviously preferable to Algorithm 5.1, as a bit of foreshadowing, we note that the large eigenvalue problem that results for this variable-distance certificate is quite different from (5.5) and has its own downsides. We now describe this certificate to complete Algorithms 6.1 and 6.2.

6.1. A 2D level-set test for variable-distance pairs. Suppose γ is both a singular value of $G(x, y)$ and $G(x, y + x\eta)$, with respective left and right singular vectors pairs u, v and \hat{u}, \hat{v} . Via Theorem 5.1 for $G(x, y)$ and a similar argument as its

Algorithm 6.2. Trisection.**Input:** A nonnormal matrix $A \in \mathbb{C}^{n \times n}$ with $\alpha(A) < 0$ and a tolerance $\gamma_{\text{tol}} > 0$.**Output:** $\text{ub}^{-1} \approx \mathcal{K}(A)$ (continuous-time).

```

1: lb  $\leftarrow$  0
2: ub  $\leftarrow$   $g(x_0, y_0)$  for some  $x_0 > 0$  and  $y_0 \in \mathbb{R}$ 
3: while (ub - lb)  $>$  ub  $\cdot$   $\gamma_{\text{tol}}$  do
4:   diff  $\leftarrow$  ub - lb
5:    $\eta \leftarrow \frac{2}{3} \cdot$  diff
6:    $\gamma \leftarrow$  lb +  $\eta$ 
7:   Perform the 2D level-set test of subsection 6.1 with current  $\gamma$  and  $\eta$ 
8:   if test finds any level-set points then
9:     ub =  $\gamma$ 
10:  else
11:    lb = lb +  $\frac{1}{3} \cdot$  diff
12:  end if
13: end while

```

NOTE: While [9, p. 358] states that their trisection-based $\tau(A, B)$ algorithm converges “to any prescribed absolute accuracy,” any desired relative accuracy can be obtained by simply choosing a stopping condition like the one we have used here in line 3.

proof for $G(x, y + x\eta)$, we have the following two Hamiltonian eigenvalue problems:

$$(6.3a) \quad \begin{bmatrix} A - xI & \gamma xI \\ -\gamma xI & xI - A^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix} = \mathbf{i}y \begin{bmatrix} v \\ u \end{bmatrix},$$

$$(6.3b) \quad \begin{bmatrix} A - x(1 + \mathbf{i}\eta)I & \gamma xI \\ -\gamma xI & x(1 - \mathbf{i}\eta)I - A^* \end{bmatrix} \begin{bmatrix} \hat{v} \\ \hat{u} \end{bmatrix} = \mathbf{i}y \begin{bmatrix} \hat{v} \\ \hat{u} \end{bmatrix}.$$

These two Hamiltonian matrices have an eigenvalue in common if

$$(6.4) \quad \begin{bmatrix} A - xI & \gamma xI \\ -\gamma xI & xI - A^* \end{bmatrix} W + W \begin{bmatrix} A^* - x(1 - \mathbf{i}\eta)I & -\gamma xI \\ \gamma xI & x(1 + \mathbf{i}\eta)I - A \end{bmatrix} = 0$$

has a nonzero solution $W \in \mathbb{C}^{2n \times 2n}$. Separating out the terms involving x , we have

$$(6.5) \quad \left(\begin{bmatrix} A & 0 \\ 0 & -A^* \end{bmatrix} W + W \begin{bmatrix} A^* & 0 \\ 0 & -A \end{bmatrix} \right) - x \left(\begin{bmatrix} I & -\gamma I \\ \gamma I & -I \end{bmatrix} W + W \begin{bmatrix} (1 - \mathbf{i}\eta)I & \gamma I \\ -\gamma I & -(1 + \mathbf{i}\eta)I \end{bmatrix} \right) = 0.$$

Rewriting both Sylvester forms using the vectorize operator, and letting $w = \text{vec}(W)$, we have the following generalized eigenvalue problem:

$$(6.6) \quad \mathcal{B}_1 w = x \mathcal{B}_2 w, \quad \text{where}$$

$$\mathcal{B}_1 = I_{2n} \otimes \begin{bmatrix} A & 0 \\ 0 & -A^* \end{bmatrix} + \begin{bmatrix} \bar{A} & 0 \\ 0 & -A^\top \end{bmatrix} \otimes I_{2n},$$

$$\mathcal{B}_2 = I_{2n} \otimes \begin{bmatrix} I & -\gamma I \\ \gamma I & -I \end{bmatrix} + \begin{bmatrix} (1 - \mathbf{i}\eta)I & -\gamma I \\ \gamma I & -(1 + \mathbf{i}\eta)I \end{bmatrix} \otimes I_{2n}.$$

To check if (6.2) holds with this variable-distance certificate, we proceed similarly to our fixed-distance certificate from subsection 5.2. Thus, after initially computing all the positive real eigenvalues of (6.6), we then apply Theorem 5.1 to each of these candidate vertical lines to see if we detect any points on a γ -level set (or lower) of $g(x, y)$.

Key Remark 6.3. We now explain how the designs of our 2D level-set tests intentionally deviate from the $\tau(A, B)$ procedure Gu proposed in [14, pp. 995–997] and how this results in much better reliability. If one were to more closely follow Gu’s procedure as it is written, for each positive real eigenvalue $\hat{x} > 0$ of (6.6), one would instead compute the eigenvalues of the two matrices in (6.3) and then check whether these two spectra have any imaginary eigenvalues in common. If a shared imaginary eigenvalue is detected, then a pair of level-set points a distance $x\eta$ apart has been detected⁴ and $\gamma \geq \mathcal{K}(A)^{-1}$ must hold. Otherwise, if none of the pairs of eigenvalue problems from (6.3) share imaginary eigenvalues, then following Gu would mean asserting that lower bound (6.2) must hold. While Gu’s procedure is sound in exact arithmetic, trying to assert whether two matrices share an (imaginary) eigenvalue is exceptionally difficult to do reliably in the presence of rounding errors. Furthermore, this matching also assumes that the real eigenvalues \hat{x} have been computed accurately enough such that the matrices (6.3) would indeed share an eigenvalue $i\hat{y}$, assuming (\hat{x}, \hat{y}) satisfies (6.1), which is another source of numerical uncertainty; see [15, section 5]. Suppose that there are indeed points a distance $x\eta$ apart vertically on the γ -level set of $g(x, y)$, but rounding errors prevent their detection. In this case, asserting that (6.2) holds may be erroneous. This can be a critical failure because the lower bound is only true *if no such pair of level-set points exists*, not if the procedure fails to detect them due to numerical problems! This is a major reason why Gu’s procedure can have such numerical difficulties, particularly when η is small. Though [15] improves Gu’s procedure to make it faster, it too follows the same the idea of checking whether or not two matrices share imaginary eigenvalues and thus also inherits these numerical problems. However, in the course of our work here, we have realized that Gu’s $\tau(A, B)$ procedure ironically follows Theorem 2.1 and Corollary 2.2 too closely, because, as it turns out, checking whether or not level-set points are a distance $x\eta$ (or η) apart is *entirely unnecessary*. While indeed (6.2) must hold if no points satisfy (6.1), another sufficient condition for (6.2) to hold is that the procedure itself does not generate *any* level-set points whatsoever, on the γ -level set or lower and/or as pairs or single points. If any level-set points are detected, clearly $\gamma \geq \mathcal{K}(A)^{-1}$ holds, while none being generated implies (6.1) cannot hold, which in turn asserts that (6.2) must hold. Note that our new way of performing these 2D level-set tests is not only a theoretical improvement; we initially designed our tests to check for common imaginary eigenvalues, like Gu’s procedure, but discovered that the aforementioned numerical issues were impeding the reliability of the codes. Applying our modifications to the $\tau(A, B)$ algorithms of [14, 9, 15] should improve their reliability as well.

6.2. Properties of the eigensystem $\mathcal{B}_1 w = x\mathcal{B}_2 w$ and its solution. Unlike (5.5), which can at least be numerically deflated to an order $2n^2$ generalized eigenvalue problem, we now show that (6.6) cannot be similarly reduced, as \mathcal{B}_2 is generally nonsingular. In fact, \mathcal{B}_2 has an explicit inverse.

⁴Note that while the bottom of [14, p. 996] says that detected points would be on the γ -level set, technically that only holds if γ is also the minimum singular value at both of these points, which may or may not be true.

THEOREM 6.4. *Let $C := \begin{bmatrix} aI & -bI \\ bI & -aI \end{bmatrix} \in \mathbb{C}^{2n \times 2n}$ with $a, b \in \mathbb{C}$ and $b \neq 0$, and let $\mathcal{M} := I_{2k} \otimes C + C \otimes I_{2k}$. For $k \in \mathbb{N}$ and $s \in \mathbb{C}$, define matrix $\mathcal{D} := sI_{4kn} + \mathcal{M}$ and scalar $\beta := s^2 + 4(b^2 - a^2)$. Then \mathcal{D} is invertible with inverse*

$$(6.7) \quad \mathcal{D}^{-1} = s^{-1}I_{4kn} - \beta^{-1}\mathcal{U}_k (I_{2kn} - 2s^{-1}I_k \otimes C) \mathcal{V}_k$$

if and only if s and β are both nonzero, where \mathcal{U}_k and \mathcal{V}_k are defined as in (5.6). Moreover, if $k = n$, then the following simpler formula holds:

$$(6.8) \quad \mathcal{D}^{-1} = s^{-1}I_{4n^2} - \beta^{-1}\mathcal{M} + (s\beta)^{-1}\mathcal{M}^2.$$

Proof. Via Lemma 5.2 applied to \mathcal{M} , we have that $\mathcal{D} = sI_{4kn} + \mathcal{U}_k\mathcal{V}_k$. As the eigenvalues of C are $\pm\sqrt{a^2 - b^2}$, by Theorem 1.3, the eigenvalues of \mathcal{M} are $\pm 2\sqrt{a^2 - b^2}$ and zero. Thus, \mathcal{D} is invertible if and only if $-s$ is not equal to any of these eigenvalues, which is equivalent to $s \neq 0$ and $\beta \neq 0$. For (6.7), we apply the Sherman–Morrison–Woodbury formula to $(sI_{4kn} + \mathcal{U}_k\mathcal{V}_k)^{-1}$ and use $\mathcal{V}_k\mathcal{U}_k = 2I_k \otimes C$ from Lemma 5.2, yielding

$$\begin{aligned} \mathcal{D}^{-1} &= s^{-1}I_{4kn} - s^{-1}\mathcal{U}_k (I_{2kn} + s^{-1}\mathcal{V}_k\mathcal{U}_k)^{-1} s^{-1}\mathcal{V}_k \\ &= s^{-1}I_{4kn} - s^{-1}\mathcal{U}_k (sI_{2kn} + 2I_k \otimes C)^{-1} \mathcal{V}_k \\ &= s^{-1}I_{4kn} - s^{-1}\mathcal{U}_k (I_k \otimes (sI_{2n} + 2C))^{-1} \mathcal{V}_k \\ &= s^{-1}I_{4kn} - s^{-1}\mathcal{U}_k (I_k \otimes (sI_{2n} + 2C))^{-1} \mathcal{V}_k \\ &= s^{-1}I_{4kn} - s^{-1}\mathcal{U}_k \left(I_k \otimes \frac{1}{s^2 + 4(b^2 - a^2)} \begin{bmatrix} (s-2a)I & 2bI \\ -2bI & (s+2a)I \end{bmatrix} \right) \mathcal{V}_k \\ &= s^{-1}I_{4kn} - s^{-1}\beta^{-1}\mathcal{U}_k (sI_{2kn} - 2I_k \otimes C) \mathcal{V}_k. \end{aligned}$$

For (6.8), now with $k = n$, first note that $C^2 = \phi I_{2n}$ with $\phi = (a^2 - b^2)$, while

$$\begin{aligned} \mathcal{M}^2 &= (I_{2n} \otimes C^2 + C^2 \otimes I_{2n}) + (C \otimes C + C \otimes C) = 2(\phi I_{4n^2} + C \otimes C), \\ \mathcal{M}^3 &= 2(\phi I_{4n^2} + C \otimes C) \mathcal{M} = 2\phi\mathcal{M} + 2(C \otimes C^2 + C^2 \otimes C) = 4\phi\mathcal{M}, \end{aligned}$$

where we have used the mixed-product property of \otimes . Then by multiplying (6.8) with $sI_{4n^2} + \mathcal{M}$ and noting that $\beta - s^2 + 4\phi = 0$, it follows that

$$\begin{aligned} I_{4n^2} + s^{-1}\mathcal{M} - s\beta^{-1}\mathcal{M} + (s\beta)^{-1}\mathcal{M}^3 &= I_{4n^2} + \left(s^{-1} - s\beta^{-1} + 4\phi (s\beta)^{-1} \right) \mathcal{M} \\ &= I_{4n^2} + (s\beta)^{-1} (\beta - s^2 + 4\phi) \mathcal{M} = I_{4n^2}. \quad \square \end{aligned}$$

Applying Theorem 6.4 to \mathcal{B}_2 , with $a := 1$, $b := \gamma \neq 0$, $s := -i\eta \neq 0$, we see that $\beta \neq 0$ holds if $\eta \neq \pm 2\sqrt{\gamma^2 - 1}$ ($\notin \mathbb{R}$ if $|\gamma| < 1$); hence \mathcal{B}_2 is generically invertible with

$$(6.9) \quad \mathcal{B}_2^{-1} = -(i\eta)^{-1}I_{4n^2} - \beta^{-1}\mathcal{B}_2 - (i\eta\beta)^{-1}\mathcal{B}_2^2,$$

where $\beta = 4(\gamma^2 - 1) - \eta^2$. On the upside, one could thus perform our variable-distance certificate by computing the eigenvalues of $\mathcal{B}_2^{-1}\mathcal{B}_1$, while our fixed-distance certificate requires solving the general eigenvalue problem (5.5). However, using $\mathcal{B}_2^{-1}\mathcal{B}_1$ can introduce numerical issues, and the order of this eigenvalue problem is $4n^2$, whereas for the fixed-distance case, we could instead solve the numerically deflated order $2n^2$ problem discussed in subsection 5.3. Furthermore, from (6.9), it is clear that $\|\mathcal{B}_2^{-1}\| \rightarrow \infty$ as $\eta \rightarrow 0$; in the presence of rounding errors, this may make it quite

difficult to reliably ascertain which eigenvalues of $\mathcal{B}_2^{-1}\mathcal{B}_1$ should be considered real-valued.

At this point, one might ask if it would have been better to consider a variable-distance certificate using a horizontal orientation instead of a vertical one. We consider this in Appendix A and note here that it results in the matrix pencil $\mathcal{B}_1 - \lambda\tilde{\mathcal{B}}_2$, where, like \mathcal{B}_2 , the matrix $\tilde{\mathcal{B}}_2$ is nonsingular but becomes singular as $\eta \rightarrow 0$.

6.3. Adapting divide-and-conquer for $\mathcal{B}_1 w = x\mathcal{B}_2 w$. We now describe the computations needed for a divide-and-conquer version of our variable-distance certificate. Note that the numerical reliability of this approach might be different than our fixed-distance certificate, as (6.6) only has finite eigenvalues, while (5.5) has finite and infinite eigenvalues. For $w \in \mathbb{C}^{4n^2}$, $\mathcal{B}_2 w$ can be efficiently obtained by storing \mathcal{B}_2 in a sparse format. Computing $\mathcal{B}_1 w$ is $\mathcal{O}(n^3)$ work, which is done by vectorizing the first matrix in (6.5), i.e., $\text{vec} \left(\begin{bmatrix} A & 0 \\ 0 & -A^* \end{bmatrix} W + W \begin{bmatrix} A^* & 0 \\ 0 & -A \end{bmatrix} \right)$, where $W \in \mathbb{C}^{2n \times 2n}$ and $w = \text{vec}(W)$. To obtain $w = (\mathcal{B}_1 - s\mathcal{B}_2)^{-1}y$ for $y \in \mathbb{C}^{4n^2}$, consider $(\mathcal{B}_1 - s\mathcal{B}_2)w = y$. This “unvectorizes” into (6.4) provided that x is replaced by s and the zero on its right-hand side is replaced by Y , where $y = \text{vec}(Y)$. Vectorizing the solution of the resulting Sylvester equation yields $w = (\mathcal{B}_1 - s\mathcal{B}_2)^{-1}y$ in $\mathcal{O}(n^3)$ work.

7. Algorithms for discrete-time $\mathcal{K}(A)$. To adapt Algorithms 5.1, 6.1, and 6.2 to compute discrete-time Kreiss constants, we need to develop discrete-time versions of the 2D level-set tests from subsection 5.2 and subsection 6.1; the other necessary components for the two optimization-with-restarts algorithms have already been discussed in subsection 3.2. We begin by developing discrete-time analogues of Theorem 4.1 and Corollary 4.2. For the discrete-time case, we now additionally assume that $0 \notin \Lambda(A)$.

THEOREM 7.1. *For $A \in \mathbb{C}^{n \times n}$ with $\rho(A) < 1$, let $\gamma \in [0, 1)$, $\eta \geq 0$, and (r_*, θ_*) be a global minimizer of (3.7). If $\mathcal{K}(A)^{-1} \leq \gamma$ and $\eta \in (0, \frac{2(r_*-1)}{1+\gamma}(\gamma - \mathcal{K}(A)^{-1})]$, then there exist an $r > 1$ and $\theta \in [0, 2\pi)$ such that*

$$(7.1) \quad h(r, \theta) = h(r + \eta, \theta) = \gamma.$$

COROLLARY 7.2. *For $A \in \mathbb{C}^{n \times n}$ with $\rho(A) < 1$, let $\gamma \in [0, 1)$, $\eta \geq 0$, and (r_*, θ_*) be a global minimizer of (3.7). If there do not exist any pairs $r, \theta \in \mathbb{R}$ with $r > 1$ such that (7.1) holds, then*

$$(7.2) \quad \mathcal{K}(A)^{-1} > \gamma - \frac{\eta(1+\gamma)}{2(r_*-1)}.$$

Proof of Theorem 7.1. If $\gamma = \mathcal{K}(A)^{-1}$, the proof is trivially satisfied with $\eta = 0$, so assume that $\gamma \in (\mathcal{K}(A)^{-1}, 1)$. Since $\rho(A) < 1$, it follows that $\lim_{r \rightarrow 1^+} h(r, \theta) = \infty$ for any $\theta \in \mathbb{R}$, and so $h(r_*, \theta_*) = \mathcal{K}(A)^{-1}$ with $r_* > 1$. Now consider the strict lower level set $\mathcal{L}_\gamma := \{(r, \theta) : h(r, \theta) < \gamma, r > 1\}$, which is clearly open and also bounded; see [27, Theorem 3.2]. Let \mathcal{L} be the (open) connected component of \mathcal{L}_γ such that $(r_*, \theta_*) \in \mathcal{L}$, and let $\mathcal{G} := \mathcal{L}^H$, i.e., the simply connected hull of \mathcal{L} .

By continuity of $h(r, \theta)$ and the boundedness of \mathcal{G} , there must exist $b_1, b_2 \in \text{bd } \mathcal{G}$,

$$b_1 = (r_* - \eta_1, \theta_*) \quad \text{and} \quad b_2 = (r_* + \eta_2, \theta_*),$$

such that

$$(7.3) \quad h(r_* - \eta_1, \theta_*) = h(r_* + \eta_2, \theta_*) = \gamma,$$

where $\eta_1, \eta_2 > 0$. Furthermore, we can assume that η_1 and η_2 are the smallest positive values such that (7.3) holds with $b_1, b_2 \in \text{bd } \mathcal{G}$. Defining

$$(7.4) \quad \tilde{h}(r) := h(r + 1, \theta) = \frac{\sigma_{\min}((r + 1)e^{i\theta}I - A)}{r},$$

$\tilde{h}(r)$ has a global minimizer $r_* - 1$ on domain $(0, \infty)$ and its numerator has a GLC of 1. Applying Lemma 4.5 to $q(r) := \tilde{h}(r)$ with $a = r_* - \eta_1 - 1$ and $b = r_* + \eta_2 - 1$ yields

$$\eta_1 + \eta_2 \geq \frac{2(r_* - 1)}{1 + \gamma}(\gamma - \mathcal{K}(A)^{-1}).$$

Now suppose that $\eta \in (0, \frac{2(r_* - 1)}{1 + \gamma}(\gamma - \mathcal{K}(A)^{-1})]$ so $\eta \leq \eta_1 + \eta_2$. If $\eta = \eta_1 + \eta_2$, (7.1) is clearly satisfied, so instead assume that $\eta < \eta_1 + \eta_2$. Considering the set $\hat{\mathcal{G}} := \{(r - \eta, \theta) : (r, \theta) \in \mathcal{G}\}$, without loss of generality, we can assume $\theta_* = 0$, and so $\hat{\mathcal{G}}$ is simply \mathcal{G} shifted left by the amount η . Thus, as in the end of the proof of Theorem 4.6, the rest of the argument follows analogously to the last paragraph of the proof of Theorem 4.1. \square

7.1. A 1D circular level-set test. For all of our discrete-time $\mathcal{K}(A)$ algorithms, we will also need the following theorem.

THEOREM 7.3. *Given $\gamma, r, \theta \in \mathbb{R}$, with $\gamma \geq 0$ and $r \neq 1$, γ is a singular value of $H(r, \theta)$ defined as in (3.8) if and only if $e^{i\theta}$ is an eigenvalue of the symplectic matrix pencil*

$$(7.5) \quad \begin{bmatrix} A & \gamma(r - 1)I \\ 0 & rI \end{bmatrix} - \lambda \begin{bmatrix} rI & 0 \\ \gamma(r - 1)I & A^* \end{bmatrix}.$$

Furthermore, if A is invertible and $r \neq 0$, zero is not an eigenvalue of (7.5), and the matrix pencil is regular.

Proof. It is easy to verify that (7.5) is symplectic, and, under the additional assumptions, also regular and that zero cannot be an eigenvalue. Now suppose γ is a singular value $H(r, \theta)$ with left and right singular vectors u and v , and so

$$\gamma \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} H(r, \theta) & 0 \\ 0 & H(r, \theta)^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix} \Leftrightarrow \gamma(r - 1) \begin{bmatrix} u \\ -e^{i\theta}v \end{bmatrix} = \begin{bmatrix} re^{i\theta}I - A & 0 \\ 0 & e^{i\theta}A^* - rI \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix}.$$

Rearranging terms, this is equivalent to

$$\gamma(r - 1) \begin{bmatrix} u \\ 0 \end{bmatrix} + \begin{bmatrix} A & 0 \\ 0 & rI \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix} = e^{i\theta} \gamma(r - 1) \begin{bmatrix} 0 \\ v \end{bmatrix} + e^{i\theta} \begin{bmatrix} rI & 0 \\ 0 & A^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix}. \quad \square$$

Similar to Theorem 5.1, note that the unimodular eigenvalues of (7.5) correspond to points that are either on the γ -level set of $h(r, \theta)$ or on lower level sets. The structure-preserving eigensolvers of [2, 3] can also be used for more reliable detection of unimodular eigenvalues of symplectic pencils like (7.5).

7.2. Adapting Algorithm 5.1 for discrete-time $\mathcal{K}(A)$. To create our first discrete-time 2D level-set test, we will again look for pairs of points a fixed distance $\eta \geq 0$ apart, but now we will do this along rays from the origin; i.e., for $\hat{\theta} \in [0, 2\pi)$, we check if $h(\hat{r}, \hat{\theta}) = h(\hat{r} + \eta, \hat{\theta}) = \gamma$ holds for some $\hat{r} > 1$. Suppose γ is a singular

value of both $H(r, \theta)$ and $H(r + \eta, \theta)$ with respective left and right singular vector pairs u, v and \hat{u}, \hat{v} . Then, by Theorem 7.3, we have that

$$(7.6a) \quad \begin{bmatrix} A & \gamma(r-1)I \\ 0 & rI \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix} = e^{i\theta} \begin{bmatrix} rI & 0 \\ \gamma(r-1)I & A^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix},$$

$$(7.6b) \quad \begin{bmatrix} A & \gamma(r+\eta-1)I \\ 0 & (r+\eta)I \end{bmatrix} \begin{bmatrix} \hat{v} \\ \hat{u} \end{bmatrix} = e^{i\theta} \begin{bmatrix} (r+\eta)I & 0 \\ \gamma(r+\eta-1)I & A^* \end{bmatrix} \begin{bmatrix} \hat{v} \\ \hat{u} \end{bmatrix},$$

which we denote $M - \lambda N$ and $\widetilde{M} - \lambda \widetilde{N}$, respectively. Now define $W = \begin{bmatrix} v \\ u \end{bmatrix} \begin{bmatrix} \hat{v}^* & \hat{u}^* \end{bmatrix} \neq 0$. Multiplying the two equations above from the right side, respectively, by $\begin{bmatrix} \hat{v}^* & \hat{u}^* \end{bmatrix} \widetilde{M}^*$ and $\begin{bmatrix} \hat{v}^* & \hat{u}^* \end{bmatrix} N^*$, yields

$$(7.7a) \quad MW\widetilde{M}^* = e^{i\theta}NW\widetilde{M}^*,$$

$$(7.7b) \quad \widetilde{M}W^*N^* = e^{i\theta}\widetilde{N}W^*N^*.$$

If we take the conjugate transpose of (7.7b) and then multiply it by $e^{i\theta}$, we obtain

$$(7.8) \quad NW\widetilde{N}^* = e^{i\theta}NW\widetilde{M}^*.$$

Subtracting (7.8) from (7.7a) yields

$$(7.9) \quad MW\widetilde{M}^* - NW\widetilde{N}^* = 0,$$

so since $M - \lambda N$ and $\widetilde{N}^* - \lambda \widetilde{M}^*$ are both regular, [10, Theorem 1] states that these two pencils must share an eigenvalue if $W \neq 0$ solves the equation above. Moreover, if this shared eigenvalue is $e^{i\theta}$, then $e^{-i\theta}$ is an eigenvalue of $\widetilde{N} - \lambda \widetilde{M}$, which in turn implies that $e^{i\theta}$ is an eigenvalue of $\widetilde{M} - \lambda \widetilde{N}$. Thus, (7.9) having a nonzero solution W is a necessary condition for the two pencils in (7.6) to have eigenvalue $e^{i\theta}$ in common.

We now want to separate out the r terms of $MW\widetilde{M}^*$ and $NW\widetilde{N}^*$. First, we have

$$(7.10) \quad \widetilde{M}^* = \begin{bmatrix} A^* & 0 \\ \gamma(r+\eta-1)I & (r+\eta)I \end{bmatrix} \quad \text{and} \quad \widetilde{N}^* = \begin{bmatrix} (r+\eta)I & \gamma(r+\eta-1)I \\ 0 & A \end{bmatrix}.$$

Then $MW\widetilde{M}^*$ is

$$(7.11) \quad \left(\begin{bmatrix} A & -\gamma I \\ 0 & 0 \end{bmatrix} + r \begin{bmatrix} 0 & \gamma I \\ 0 & I \end{bmatrix} \right) W \left(\begin{bmatrix} A^* & 0 \\ \gamma(\eta-1)I & \eta I \end{bmatrix} + r \begin{bmatrix} 0 & 0 \\ \gamma I & I \end{bmatrix} \right),$$

which is equal to

$$(7.12) \quad \begin{bmatrix} A & -\gamma I \\ 0 & 0 \end{bmatrix} W \begin{bmatrix} A^* & 0 \\ \gamma(\eta-1)I & \eta I \end{bmatrix} + r \left(\begin{bmatrix} A & -\gamma I \\ 0 & 0 \end{bmatrix} W \begin{bmatrix} 0 & \gamma I \\ \gamma I & I \end{bmatrix} + \begin{bmatrix} 0 & \gamma I \\ 0 & I \end{bmatrix} W \begin{bmatrix} A^* & 0 \\ \gamma(\eta-1)I & \eta I \end{bmatrix} \right) + r^2 \begin{bmatrix} 0 & \gamma I \\ 0 & I \end{bmatrix} W \begin{bmatrix} 0 & 0 \\ \gamma I & I \end{bmatrix}.$$

Vectorizing the above equation, with $w = \text{vec}(W)$, yields

$$(7.13) \quad \begin{bmatrix} A & \gamma(\eta-1)I \\ 0 & \eta I \end{bmatrix} \otimes \begin{bmatrix} A & -\gamma I \\ 0 & 0 \end{bmatrix} w + r \left(\begin{bmatrix} 0 & \gamma I \\ 0 & I \end{bmatrix} \otimes \begin{bmatrix} A & -\gamma I \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} A & \gamma(\eta-1)I \\ 0 & \eta I \end{bmatrix} \otimes \begin{bmatrix} 0 & \gamma I \\ 0 & I \end{bmatrix} \right) w + r^2 \begin{bmatrix} 0 & \gamma I \\ 0 & I \end{bmatrix} \otimes \begin{bmatrix} 0 & \gamma I \\ 0 & I \end{bmatrix} w,$$

which we will abbreviate as

$$(7.14) \quad \mathcal{M}_0 w + r \mathcal{M}_1 w + r^2 \mathcal{M}_2 w.$$

Likewise, $NW\tilde{N}^*$ is

$$(7.15) \quad \left(\begin{bmatrix} 0 & 0 \\ -\gamma I & A^* \end{bmatrix} + r \begin{bmatrix} I & 0 \\ \gamma I & 0 \end{bmatrix} \right) W \left(\begin{bmatrix} \eta I & \gamma(\eta-1)I \\ 0 & A \end{bmatrix} + r \begin{bmatrix} I & \gamma I \\ 0 & 0 \end{bmatrix} \right),$$

which is equal to

$$(7.16) \quad \begin{bmatrix} 0 & 0 \\ -\gamma I & A^* \end{bmatrix} W \begin{bmatrix} \eta I & \gamma(\eta-1)I \\ 0 & A \end{bmatrix} \\ + r \left(\begin{bmatrix} 0 & 0 \\ -\gamma I & A^* \end{bmatrix} W \begin{bmatrix} I & \gamma I \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} I & 0 \\ \gamma I & 0 \end{bmatrix} W \begin{bmatrix} \eta I & \gamma(\eta-1)I \\ 0 & A \end{bmatrix} \right) + r^2 \begin{bmatrix} I & 0 \\ \gamma I & 0 \end{bmatrix} W \begin{bmatrix} I & \gamma I \\ 0 & 0 \end{bmatrix}.$$

Similarly vectorizing this gives

$$(7.17) \quad \begin{bmatrix} \eta I & 0 \\ \gamma(\eta-1)I & A^\top \end{bmatrix} \otimes \begin{bmatrix} 0 & 0 \\ -\gamma I & A^* \end{bmatrix} w \\ + r \left(\begin{bmatrix} I & 0 \\ \gamma I & 0 \end{bmatrix} \otimes \begin{bmatrix} 0 & 0 \\ -\gamma I & A^* \end{bmatrix} + \begin{bmatrix} \eta I & 0 \\ \gamma(\eta-1)I & A^\top \end{bmatrix} \otimes \begin{bmatrix} I & 0 \\ \gamma I & 0 \end{bmatrix} \right) w + r^2 \begin{bmatrix} I & 0 \\ \gamma I & 0 \end{bmatrix} \otimes \begin{bmatrix} I & 0 \\ \gamma I & 0 \end{bmatrix} w,$$

which we will abbreviate as

$$(7.18) \quad \mathcal{N}_0 w + r \mathcal{N}_1 w + r^2 \mathcal{N}_2 w.$$

Therefore, we finally have the following quadratic eigenvalue problem:

$$(7.19) \quad (\mathcal{M}_0 - \mathcal{N}_0) w + r (\mathcal{M}_1 - \mathcal{N}_1) w + r^2 (\mathcal{M}_2 - \mathcal{N}_2) w = 0,$$

which we will abbreviate as

$$(7.20) \quad \mathcal{Q}_0 w + r \mathcal{Q}_1 w + r^2 \mathcal{Q}_2 w = 0.$$

Thus, to perform our fixed-distance discrete-time $\mathcal{K}(A)$ certificate for adapting Algorithm 5.1, we first compute all the real eigenvalues $r > 1$ of (7.20), which specify a set of concentric circles centered at the origin on which we may find level set points. Then, for each of these candidate radii, we apply Theorem 7.3 to see if we indeed detect any points on a γ -level set (or lower) of $h(r, \theta)$. Like our continuous-time algorithms, this procedure is also $\mathcal{O}(n^6)$ work (albeit with a larger constant term) when using dense (quadratic) eigensolvers.

Note that for a generic quadratic eigenvalue problem of the form (7.20), if either \mathcal{Q}_0 or \mathcal{Q}_2 is nonsingular, then the problem is well posed; i.e., it has at least one solution and not infinitely many; see [1, p. 283]. While for our quadratic eigenvalue problem \mathcal{Q}_2 is singular (see Lemma B.4 in Appendix B), we now show that \mathcal{Q}_0 is generically nonsingular, and thus (7.20) is guaranteed to be well posed.

THEOREM 7.4. *Let $\gamma, \eta \in \mathbb{R}$ both be positive. Then \mathcal{Q}_0 from (7.20) is nonsingular if and only if A is nonsingular and γ is not a singular value of A .*

Proof. For any vector $w \in \mathbb{C}^{4n^2}$, suppose $\mathcal{Q}_0 w = 0$, i.e.,

$$\left(\begin{bmatrix} \bar{A} & \gamma(\eta-1)I \\ 0 & \eta I \end{bmatrix} \otimes \begin{bmatrix} A & -\gamma I \\ 0 & 0 \end{bmatrix} \right) w - \left(\begin{bmatrix} \eta I & 0 \\ \gamma(\eta-1)I & A^\top \end{bmatrix} \otimes \begin{bmatrix} 0 & 0 \\ -\gamma I & A^* \end{bmatrix} \right) w = 0,$$

which holds if and only if

$$(7.21) \quad \begin{bmatrix} A & -\gamma I \\ 0 & 0 \end{bmatrix} W \begin{bmatrix} A^* & 0 \\ \gamma(\eta - 1)I & \eta I \end{bmatrix} - \begin{bmatrix} 0 & 0 \\ -\gamma I & A^* \end{bmatrix} W \begin{bmatrix} \eta I & \gamma(\eta - 1)I \\ 0 & A \end{bmatrix} = 0,$$

where $w = \text{vec}(W)$. By [10, Theorem 1], the generalized Sylvester equation above has a unique solution if and only if the two matrix pencils

$$\begin{bmatrix} A & -\gamma I \\ 0 & 0 \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 \\ -\gamma I & A^* \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} \eta I & \gamma(\eta - 1)I \\ 0 & A \end{bmatrix} - \lambda \begin{bmatrix} A^* & 0 \\ \gamma(\eta - 1)I & \eta I \end{bmatrix}$$

are both regular and have no eigenvalues in common. Clearly $W = 0$ satisfies (7.21), so as long as zero is the only solution, \mathcal{Q}_0 is nonsingular for the assumptions. To prove the forward direction, we thus show that these three conditions hold.

We begin with the first pencil, which is regular if, for at least one value of λ , the resulting matrix is nonsingular. Using $\lambda = -1$ results in $\begin{bmatrix} A & -\gamma I \\ -\gamma I & A^* \end{bmatrix}$. Since A is invertible, this matrix is invertible if and only if its Schur complement with respect to A is, i.e.,

$$0 \neq \det(A^* - \gamma^2 A^{-1}) = \det(AA^* - \gamma^2 I) / \det(A),$$

which by definition of singular values, holds if and only if γ is not a singular value of A . Thus, the first pencil is regular. For the second pencil, choosing $\lambda = 0$ results in matrix $\begin{bmatrix} \eta I & \gamma(\eta - 1)I \\ 0 & A \end{bmatrix}$, which is invertible since $\eta \neq 0$ and A is invertible. Thus the second pencil is also regular. Furthermore, this argument also establishes that zero cannot be an eigenvalue of the second pencil.

We now show that the two pencils do not have any eigenvalues in common. We begin by noting that some of the eigenvalues of the first pencil are infinity, while all the eigenvalues of the second pencil are finite. Let $\lambda \neq 0$ be a finite eigenvalue of the first pencil with eigenvector $\begin{bmatrix} v \\ u \end{bmatrix} \neq 0$; hence

$$(7.22) \quad \begin{bmatrix} A & -\gamma I \\ \lambda\gamma I & -\lambda A^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix} = 0 \quad \iff \quad \begin{aligned} Av &= \gamma u \\ A^*u &= \gamma v. \end{aligned}$$

Since $\gamma \neq 0$ and A is invertible, $u = 0$ if and only if $v = 0$; hence neither is zero. This means they can be rescaled to each have unit norm, and so γ is a singular value of A , a contradiction; hence $\lambda = 0$ must hold. As zero cannot be an eigenvalue of the second pencil, this part of the proof is complete.

For the reverse direction, first suppose that A is singular. Then zero must be an eigenvalue of the second pencil; hence the two pencils share zero as an eigenvalue and so \mathcal{Q}_0 is singular. Now suppose that γ is a singular value of A with left and right singular vectors \tilde{u} and \tilde{v} , and again consider (7.22). Since $\begin{bmatrix} v \\ u \end{bmatrix} := \begin{bmatrix} \tilde{v} \\ \tilde{u} \end{bmatrix} \neq 0$ is in the nullspace of the matrix given in (7.22) for any $\lambda \in \mathbb{C}$, the first pencil matrix is not regular, and so \mathcal{Q}_0 is singular. \square

Remark 7.5. It is also possible to derive a globality certificate based on arcs instead of radial segments, where for $\hat{r} > 1$, angle(s) $\hat{\theta}$ satisfying $h(\hat{r}, \hat{\theta}) = h(\hat{r}, \hat{\theta} + \eta) = \gamma$ are sought. We also considered this, but it resulted in a quadratic eigenvalue problem where both of the corresponding \mathcal{Q}_0 and \mathcal{Q}_2 matrices were singular, and so it was unclear if this alternative quadratic eigenvalue problem was well posed or not.

7.3. Adapting Algorithms 6.1 and 6.2 for discrete-time $\mathcal{K}(A)$. Now, following the variable-distance certificate idea from section 6, we derive a new discrete-time version to extend Algorithms 6.1 and 6.2 so that they can compute discrete-time Kreiss constants.

THEOREM 7.6. For $A \in \mathbb{C}^{n \times n}$ with $\rho(A) < 1$, let $\gamma \in [0, 1)$, $\eta \geq 0$, and (r_*, θ_*) be a global minimizer of (3.7). If $\mathcal{K}(A)^{-1} \leq \gamma$ and $\eta \in [0, 2(\gamma - \mathcal{K}(A)^{-1})]$, then there exist an $r > 1$ and $\theta \in [0, 2\pi)$ such that

$$(7.23) \quad h(r, \theta) = h(\beta r + \delta, \theta) = \gamma,$$

where $\beta := 1 - \delta$ and $\delta := -\frac{\eta}{1+\gamma}$.

COROLLARY 7.7. For $A \in \mathbb{C}^{n \times n}$ with $\rho(A) < 1$, let $\gamma \in [0, 1)$, $\eta \geq 0$, and (r_*, θ_*) be a global minimizer of (3.7). If there do not exist any pairs $r, \theta \in \mathbb{R}$ with $r > 1$ such that (7.23) holds, then

$$(7.24) \quad \mathcal{K}(A)^{-1} > \gamma - \frac{\eta}{2x_*}.$$

Proof of Theorem 7.6. The proof follows similarly to the proof of Theorem 7.1, except that (7.23) corresponds to a distance of $\tilde{\eta} := (\beta r + \delta) - r = \frac{\eta(r-1)}{1+\gamma}$ between the two level-set points, which leads to cancellation, and so $\eta \in [0, 2(\gamma - \mathcal{K}(A)^{-1})]$. \square

For brevity, we skip showing the lengthy derivation of the resulting discrete-time variable-distance certificate (it follows similarly to subsection 7.2) and instead just give the key parts necessary to perform the computation. Suppose γ is a singular value of both $H(r, \theta)$ and $H(\beta r + \delta, y)$ with respective left and right singular vectors pairs u, v and \hat{u}, \hat{v} . Applying Theorem 7.3 to $H(r, \theta)$ and $H(\beta r + \delta, \theta)$ yields (7.6a) and (7.6b) but now with r replaced by $\beta r + \delta$. For this modified pair of symplectic eigenvalue problems, the certificate derivation results in the large quadratic eigenvalue problem

$$(7.25) \quad \tilde{\mathcal{Q}}_0 w + r \tilde{\mathcal{Q}}_1 w + r^2 \tilde{\mathcal{Q}}_2 w = 0, \quad \text{where}$$

$$\tilde{\mathcal{Q}}_0 = \mathcal{Q}_0 \text{ given in (7.20) but with } \eta \text{ replaced by } \delta,$$

$$\tilde{\mathcal{Q}}_1 = \beta \left(\begin{bmatrix} 0 & \gamma I \\ 0 & I \end{bmatrix} \otimes \begin{bmatrix} A & -\gamma I \\ 0 & 0 \end{bmatrix} - \begin{bmatrix} I & 0 \\ \gamma I & 0 \end{bmatrix} \otimes \begin{bmatrix} 0 & 0 \\ -\gamma I & A^* \end{bmatrix} \right)$$

$$\quad + \left(\begin{bmatrix} \bar{A} & \gamma(\delta-1)I \\ 0 & \delta I \end{bmatrix} \otimes \begin{bmatrix} 0 & \gamma I \\ 0 & I \end{bmatrix} - \begin{bmatrix} \delta I & 0 \\ \gamma(\delta-1)I & A^T \end{bmatrix} \otimes \begin{bmatrix} I & 0 \\ \gamma I & 0 \end{bmatrix} \right),$$

$$\tilde{\mathcal{Q}}_2 = \beta \mathcal{Q}_2 \text{ given in (7.20).}$$

Although $\tilde{\mathcal{Q}}_2$ is singular, $\tilde{\mathcal{Q}}_0$ is nonsingular for the assumptions of Theorem 7.4, since $\eta := \delta \neq 0$. Thus, (7.25) is well posed under the same assumptions.

7.4. Adapting divide-and-conquer for discrete-time $\mathcal{K}(A)$. We now show how the real eigenvalues of (7.20) and (7.25) may be computed using divide-and-conquer. We begin with (7.20) and form its (companion) linearization

$$(7.26) \quad \begin{bmatrix} \mathcal{Q}_1 & \mathcal{Q}_0 \\ -I & 0 \end{bmatrix} z = r \begin{bmatrix} -\mathcal{Q}_2 & 0 \\ 0 & -I \end{bmatrix} z,$$

where $z = \begin{bmatrix} r w \\ w \end{bmatrix}$. Assuming an a priori upper bound $D > 0$ is known for all real eigenvalues of (7.20), divide-and-conquer will sweep the interval $[1, D]$ to find all the real-valued eigenvalues in this range. We now detail how the necessary operations with the matrices in (7.26) can all be done in at most $\mathcal{O}(n^3)$ work.

Consider doing matrix-vector products with either matrix in (7.26) and a vector $w = \begin{bmatrix} w_1 \\ w_2 \end{bmatrix}$, where $w_1, w_2 \in \mathbb{C}^{2n^2}$. The nontrivial parts of these products are

$$(7.27a) \quad \mathcal{Q}_0 w_2 = \mathcal{M}_0 w_2 - \mathcal{N}_0 w_2,$$

$$(7.27b) \quad \mathcal{Q}_1 w_1 = \mathcal{M}_1 w_1 - \mathcal{N}_1 w_1,$$

$$(7.27c) \quad \mathcal{Q}_2 w_1 = \mathcal{M}_2 w_1 - \mathcal{N}_2 w_1,$$

which are equal to the following respective vectorizations:

$$(7.28a) \quad \mathcal{Q}_0 w_2 = \text{vec} \left(\begin{bmatrix} A & -\gamma I \\ 0 & 0 \end{bmatrix} W_2 \begin{bmatrix} \gamma(\eta-1)I & 0 \\ \eta I & A \end{bmatrix} - \begin{bmatrix} 0 & 0 \\ -\gamma I & A^* \end{bmatrix} W_2 \begin{bmatrix} \eta I & \gamma(\eta-1)I \\ 0 & A \end{bmatrix} \right),$$

$$(7.28b) \quad \mathcal{Q}_1 w_1 = \text{vec} \left(\begin{bmatrix} A & -\gamma I \\ 0 & 0 \end{bmatrix} W_1 \begin{bmatrix} 0 & 0 \\ \gamma I & I \end{bmatrix} + \begin{bmatrix} 0 & \gamma I \\ -\gamma I & A^* \end{bmatrix} W_1 \begin{bmatrix} \gamma(\eta-1)I & 0 \\ 0 & \eta I \end{bmatrix} \right. \\ \left. - \begin{bmatrix} 0 & 0 \\ -\gamma I & A^* \end{bmatrix} W_1 \begin{bmatrix} I & \gamma I \\ 0 & 0 \end{bmatrix} - \begin{bmatrix} I & 0 \\ \gamma I & 0 \end{bmatrix} W_1 \begin{bmatrix} \eta I & \gamma(\eta-1)I \\ 0 & A \end{bmatrix} \right),$$

$$(7.28c) \quad \mathcal{Q}_2 w_1 = \text{vec} \left(\begin{bmatrix} 0 & I \\ 0 & \gamma I \end{bmatrix} W_1 \begin{bmatrix} 0 & 0 \\ I & \gamma I \end{bmatrix} - \begin{bmatrix} \gamma I & 0 \\ I & 0 \end{bmatrix} W_1 \begin{bmatrix} \gamma I & I \\ 0 & 0 \end{bmatrix} \right),$$

where $w_1 = \text{vec}(W_1)$ and $w_2 = \text{vec}(W_2)$. The first two of these can be obtained in $\mathcal{O}(n^3)$ work since they only involve matrix-matrix products with $2n \times 2n$ matrices. The third can be obtained in $\mathcal{O}(n^2)$ work as the number of nonzero entries in \mathcal{Q}_2 is simply $8n^2$; hence one should just store \mathcal{Q}_2 in a sparse format. This is also fortunate as when applying shift-and-invert to a generalized eigenvalue problem $Ax = \lambda Bx$, solvers such as `eigs` in MATLAB require that B is provided explicitly, even when the operator $(A - sB)^{-1}$ is given implicitly as a function handle.

Given a shift $s \in \mathbb{C}$ and a vector $y = \begin{bmatrix} y_1 \\ y_2 \end{bmatrix}$, where $y_1, y_2 \in \mathbb{C}^{2n^2}$, we now focus on

$$(7.29) \quad w = \left(\begin{bmatrix} \mathcal{Q}_1 & \mathcal{Q}_0 \\ -I & 0 \end{bmatrix} - s \begin{bmatrix} -\mathcal{Q}_2 & 0 \\ 0 & -I \end{bmatrix} \right)^{-1} y,$$

which is equivalent to

$$(7.30) \quad \left(\begin{bmatrix} \mathcal{Q}_1 & \mathcal{Q}_0 \\ -I & 0 \end{bmatrix} - s \begin{bmatrix} -\mathcal{Q}_2 & 0 \\ 0 & -I \end{bmatrix} \right) \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} \mathcal{Q}_1 + s\mathcal{Q}_2 & \mathcal{Q}_0 \\ -I & sI \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \end{bmatrix}.$$

The bottom block row provides

$$(7.31) \quad s w_2 = y_2 + w_1.$$

Multiplying the top block row of (7.30) by s and then substituting in (7.31), we get

$$(7.32) \quad \mathcal{Q}_0 w_1 + s\mathcal{Q}_1 w_1 + s^2 \mathcal{Q}_2 w_1 = s y_1 - \mathcal{Q}_0 y_2 =: \hat{y}.$$

Using (7.28) to obtain the four matrix-vector products above, we can then solve for w_1 via solving the following generalized continuous-time algebraic Sylvester equation:

$$(7.33) \quad M W_1 \tilde{M}^* - N W_1 \tilde{N}^* = \hat{Y},$$

where $w_1 = \text{vec}(W_1)$, $\hat{y} = \text{vec}(\hat{Y})$, and the matrix pairs M, N and \tilde{M}^*, \tilde{N}^* are respectively given in (7.6a) and (7.10), all with r replaced with s . Per [17], solving (7.33) can be done in $\mathcal{O}(n^3)$ work. Finally, w_2 is obtained via (7.31).

For (7.25), only a few minor modifications to the divide-and-conquer variant we have just explained for (7.20) are necessary. As $\tilde{\mathcal{Q}}_0$ is equal to \mathcal{Q}_0 with $\eta := \delta$ and $\tilde{\mathcal{Q}}_2 = \beta \mathcal{Q}_2$, the first and third equations in (7.28) can be used to do the corresponding matrix-vector products. For $w_1 \in \mathbb{C}^{2n^2}$, we have

$$(7.34) \quad \tilde{\mathcal{Q}}_1 w_1 = \text{vec} \left(\beta \left(\begin{bmatrix} A & -\gamma I \\ 0 & 0 \end{bmatrix} W_1 \begin{bmatrix} 0 & 0 \\ \gamma I & I \end{bmatrix} - \begin{bmatrix} 0 & 0 \\ -\gamma I & A^* \end{bmatrix} W_1 \begin{bmatrix} I & \gamma I \\ 0 & 0 \end{bmatrix} \right) \right. \\ \left. + \left(\begin{bmatrix} 0 & \gamma I \\ 0 & I \end{bmatrix} W_1 \begin{bmatrix} \gamma(\delta-1)I & 0 \\ \delta I & \delta I \end{bmatrix} - \begin{bmatrix} I & 0 \\ \gamma I & 0 \end{bmatrix} W_1 \begin{bmatrix} \eta I & \gamma(\delta-1)I \\ 0 & A \end{bmatrix} \right) \right).$$




FIG. 1. The solid curves plot $\frac{\alpha_\varepsilon(A)}{\varepsilon}$ (left) and $\frac{\rho_\varepsilon(A)-1}{\varepsilon}$ (right) as they vary with ε for our continuous- and discrete-time examples, respectively. The dashed lines depict the values of $\mathcal{K}(A)$ computed by our new methods, verifying their global convergence. The plots were produced using Chebfun [11] and the methods of [4].

8. Numerical experiments. To validate our methods for computing Kreiss constants, we implemented proof-of-concepts in MATLAB. The supplementary material includes the code, test examples, and a detailed description of the setup in order to reproduce the experiments in this paper. We plan to add “production-ready” implementations of our methods to a future release of ROSTAPACK [26].

We considered two 10×10 stable matrices based on demos in EigTool [31], one for continuous-time $\mathcal{K}(A)$ and a second for discrete-time $\mathcal{K}(A)$. For the former, we used $A = B - \kappa I$, where $B = \text{companion_demo}(10)$ and $\kappa = 1.001\alpha(B)$; this matrix has a large Kreiss constant, and as shown in Figure 1, (1.9) has two local maximizers for this example. For the latter, we chose $A = \frac{1}{13}B + \frac{11}{10}I$, where $B = \text{convdiff_demo}(11)$; while this matrix has a small Kreiss constant, it is interesting for testing, as, for this example, $h(r, \theta)$ has several local minimizers (see the supplementary material) and, per Figure 1, (1.4) appears to be nonsmooth (though not at the maximizer). To verify restarting, we intentionally chose starting points such that global minimizers of $g(x, y)$ and $h(r, \theta)$ would not be found in the first round of optimization. In Table 1, we provide detailed metrics on our three algorithms and see in practice that the optimization-based Algorithms 5.1 and 6.1 are much faster and more accurate than the trisection-based Algorithm 6.2. The much higher numerical accuracy of our two optimization methods is also verified by comparing to Chebfun, with Chebfun being much slower, even though it was given a small interval containing a global maximizer attaining $\mathcal{K}(A)$. In fact, Chebfun is also much slower than Algorithm 6.2, which is the slowest of our three new methods, due to the linear convergence of trisection. As can be seen, Algorithm 6.1 appears to be the best of the algorithms. See the supplementary material for additional figures showing how Algorithms 5.1 and 6.1 progress from one minimizer to the next to converge to $\mathcal{K}(A)$.

We also evaluated our theoretically faster divide-and-conquer approaches. The supplementary material includes a demo for reproducing one key experiment of the many we performed. This demo uses the computations described in subsection 5.4 and subsection 6.3 to respectively compute eigenvalues of (5.5) and (6.6) closest to a given shift, where γ and η were set to the values used in the first 2D level-set test computed by Algorithm 5.1 when computing $\mathcal{K}(A)$ of our continuous-time example.

TABLE 1

The upper half of the table shows data for our continuous-time example, while the lower half is for the discrete-time example. A dash indicates that the column is not relevant for the given method. Chebfun is simply taking the max of the chebfuns produced to make Figure 1, where Chebfun was supplied a reasonably small interval already known to contain the maximizer attaining $\mathcal{K}(A)$.

	$\mathcal{K}(A)$	z_0	Level-set tests	Restarts	Seconds
Chebfun	$1.29186707005845 \times 10^5$	—	—	—	97.28
Algorithm 5.1	$1.29186707011257 \times 10^5$	$6 + 6i$	14	2	1.95
Algorithm 6.1	$1.29186707015035 \times 10^5$	$6 + 6i$	2	2	0.76
Algorithm 6.2	$1.29181067678395 \times 10^5$	1	89	—	21.48
Chebfun	1.89501339090609	—	—	—	74.04
Algorithm 5.1	1.89501339090580	$-1 + 1i$	15	3	7.18
Algorithm 6.1	1.89501339090580	$-1 + 1i$	3	3	1.62
Algorithm 6.2	1.89501305930067	2	81	—	41.04

However, we observed that these shift-and-invert closest eigenvalue computations were somewhat unreliable when compared to using `eig` in MATLAB on (5.5) and (6.6) directly. For space reasons, we forgo many of the details and instead make a few key observations. First, sometimes the eigenvalues computed by the shift-and-invert approach did not have even a single digit of agreement to those computed by `eig`, but as these eigenvalues were often very close to the origin, this is not so surprising. Second, we also observed that `eigs` in MATLAB would sometimes fail to return one of the closest eigenvalues to a given shift; interestingly, this phenomenon can even be observed in the paper where divide-and-conquer for $\tau(A, B)$ was proposed [15, top left plot of Figure 3.1]. Third, we also observed that the large eigenvalue problems sometimes have very close conjugate pair eigenvalues, which may explain some of the aforementioned difficulties with extracting eigenvalues on the real axis accurately. As divide-and-conquer is built on the assumption that shift-and-invert eigensolvers are reliable, addressing the numerical issues above may just be a matter of picking the right sparse eigensolver. That being said, our initial testing of divide-and-conquer and its numerical reliability is merely a starting point for future work.

9. Concluding remarks. We have presented the first globally convergent algorithms for computing continuous- and discrete-time Kreiss constants to arbitrary accuracy. Algorithm 6.1, which uses our variable-distance 2D level-set tests, appears to be the fastest of the three different iterations proposed here, since it requires performing the least number of these tests. However, for continuous-time Kreiss constants, upcoming fast deflation methods [18, Chapter 3.1] may enable our fixed-distance 2D level-set test of subsection 5.2 to be computed with significantly less work, which could make Algorithm 5.1 more competitive and perhaps even faster than Algorithm 6.1. While we have outlined the key theoretical differences between the various large eigenvalue problems underlying our fixed- and variable-distance certificates, a systematic numerical evaluation could help illuminate whether there are any notable differences in terms of numerical reliability when using dense eigensolvers. Similarly, further investigation into the numerics of our asymptotically faster divide-and-conquer variants using sparse shift-and-invert eigensolvers could prove useful.

Appendix A. A 2D level-set test for variable-distance pairs with horizontal orientation. If we modify Theorem 4.6 to consider horizontal pairs of points a variable distance $\frac{x\eta}{1+\gamma}$ apart (instead of a fixed distance η apart), to induce cancellation in the proof, we obtain the following results.

THEOREM A.1. For $A \in \mathbb{C}^{n \times n}$ with $\alpha(A) < 0$, let $\gamma \in [0, 1)$, $\eta \geq 0$, and (x_*, y_*) be a global minimizer of (3.1). If $\mathcal{K}(A)^{-1} \leq \gamma$ and $\eta \in [0, 2(\gamma - \mathcal{K}(A)^{-1})]$, then there exists a pair $x, y \in \mathbb{R}$ with $x > 0$ such that

$$(A.1) \quad g(x, y) = g(\beta x, y) = \gamma,$$

where $\beta := 1 + \frac{\eta}{1+\gamma}$.

COROLLARY A.2. For $A \in \mathbb{C}^{n \times n}$ with $\alpha(A) < 0$, let $\gamma \in [0, 1)$, $\eta \geq 0$, and (x_*, y_*) be a global minimizer of (3.1). If there do not exist any pairs $x, y \in \mathbb{R}$ with $x > 0$ such that (A.1) holds, then

$$(A.2) \quad \mathcal{K}(A)^{-1} > \gamma - \frac{\eta}{2x_*}.$$

Proof of Theorem A.1. The proof follows similarly to the proof of Theorem 4.6, except that (A.1) corresponds to a distance of $\tilde{\eta} := \beta x - x = \frac{\eta x}{1+\gamma}$ between the two level-set points, which leads to cancellation, and so $\eta \in [0, 2(\gamma - \mathcal{K}(A)^{-1})]$. \square

The derivation of the corresponding verification procedure, to enable versions of Algorithms 6.1 and 6.2 using (A.1), is as follows. Suppose γ is a singular value of both $G(x, y)$ and $G(\beta x, y)$ with respective left and right singular vectors pairs u, v and \hat{u}, \hat{v} . Applying Theorem 5.1 to $G(x, y)$ and $G(\beta x, y)$ yields the following two Hamiltonian standard eigenvalue problems:

$$(A.3) \quad \begin{bmatrix} A - xI & \gamma xI \\ -\gamma xI & xI - A^* \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix} = \mathbf{i}y \begin{bmatrix} v \\ u \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} A - \beta xI & \gamma \beta xI \\ -\gamma \beta xI & \beta xI - A^* \end{bmatrix} \begin{bmatrix} \hat{v} \\ \hat{u} \end{bmatrix} = \mathbf{i}y \begin{bmatrix} \hat{v} \\ \hat{u} \end{bmatrix},$$

which have a common eigenvalue if

$$(A.4) \quad \begin{bmatrix} A - xI & \gamma xI \\ -\gamma xI & xI - A^* \end{bmatrix} W + W \begin{bmatrix} A^* - \beta xI & -\gamma \beta xI \\ \gamma \beta xI & \beta xI - A \end{bmatrix} = 0$$

has a nonzero solution $W \in \mathbb{C}^{2n \times 2n}$. Separating this into two Sylvester forms to isolate x and then vectorizing yields the generalized eigenvalue problem

$$(A.5) \quad \mathcal{B}_1 w = x \tilde{\mathcal{B}}_2 w, \quad \text{where} \quad \tilde{\mathcal{B}}_2 = I_{2n} \otimes \begin{bmatrix} I & -\gamma I \\ \gamma I & -I \end{bmatrix} + \beta \begin{bmatrix} I & -\gamma I \\ \gamma I & -I \end{bmatrix} \otimes I_{2n},$$

\mathcal{B}_1 is defined in (6.6), and $w = \text{vec}(W)$. Applying Theorem 1.3 to $\tilde{\mathcal{B}}_2$, we see that its eigenvalues are the pairwise sums of $\pm\sqrt{1 - \gamma^2}$ and $\pm\beta\sqrt{1 - \gamma^2}$; hence it is nonsingular if and only if $\beta \neq \pm 1$ and $\gamma \neq \pm 1$. Furthermore, via Lemma B.3,

$$(A.6) \quad \tilde{\mathcal{B}}_2^{-1} = 2\omega\phi\tilde{\mathcal{B}}_2 - \omega\tilde{\mathcal{B}}_2^3,$$

where $\phi := (1 + \beta^2)(1 - \gamma^2)$ and $\omega := (1 - \beta^2)^{-2}(1 - \gamma^2)^{-2}$. However, as $\eta \rightarrow 0$, $\beta \rightarrow 1$, and so $\tilde{\mathcal{B}}_2$ will become closer and closer to being singular as the algorithms progress. This suggests that it would be better numerically to solve (A.5) as a generalized eigenvalue problem, rather than solving the standard eigenvalue problem $\tilde{\mathcal{B}}_2^{-1}\mathcal{B}_1$. In terms of the number of arithmetic operations, if A is real, using divide-and-conquer to compute the positive real eigenvalues of (A.5) only involves matrix problems with real coefficients. The two eigenvalue problems in (A.3) also have real coefficients when A is real. Given $s \in \mathbb{R}$, $y = (\mathcal{B}_1 - s\tilde{\mathcal{B}}_2)^{-1}w$ can be computed by solving the Sylvester equation (A.4) with x replaced by s and its zero right-hand side replaced by Y , where $w = \text{vec}(W)$ and $y = \text{vec}(Y)$. In contrast, the computations for the divide-and-conquer variant from subsection 6.3 require complex arithmetic.

Appendix B. Some technical results.

LEMMA B.1. *Let $\{\gamma_j\}$ be the iterates produced by trisection with $\{\gamma_j\} \rightarrow \gamma_*$, and suppose that trisection terminates at the k th iterate with $|\gamma_* - \gamma_k| \leq \psi\gamma_*$ holding for some given relative error tolerance $\psi > 0$. Then $\eta_k \leq (1 + \psi)\gamma_*$.*

Proof. By construction, $\gamma_k = L + \eta_k$ and $L \geq 0$, and so $\eta_k \leq \gamma_k$ always holds. If $\gamma_k \geq \gamma_*$, then $\psi\gamma_* \geq |\gamma_* - \gamma_k| = \gamma_k - \gamma_*$, which implies $\eta_k \leq (1 + \psi)\gamma_*$. Otherwise, $\gamma_k < \gamma_*$ must hold, and so it follows that $\psi\gamma_* \geq |\gamma_* - \gamma_k| = \gamma_* - \gamma_k > \gamma_k - \gamma_*$. \square

COROLLARY B.2. *Given a fixed $\psi > 0$, if $\eta_k > (1 + \psi)\gamma_*$, then the relative error of trisection at the k th iterate is bounded below by ψ , specifically $|\gamma_* - \gamma_k| > \psi\gamma_*$.*

LEMMA B.3. *Let $C := \begin{bmatrix} aI & -bI \\ bI & -aI \end{bmatrix} \in \mathbb{C}^{2n \times 2n}$ with $a, b \in \mathbb{C}$, and let $\beta \in \mathbb{C}$. Then the matrix $\mathcal{D} := I_{2n} \otimes C + \beta C \otimes I_{2n}$ is invertible if and only if $a^2 \neq b^2$ and $\beta^2 \neq 1$, where*

$$(B.1) \quad \mathcal{D}^{-1} = 2\omega\phi D - \omega\mathcal{D}^3,$$

$$\phi := (1 + \beta^2)(a^2 - b^2), \text{ and } \omega := (1 - \beta^2)^{-2}(a^2 - b^2)^{-2}.$$

Proof. First note that $C^2 = (a^2 - b^2)I_{2n}$, while using the mixed-product property of \otimes yields

$$\begin{aligned} \mathcal{D}^2 &= (I_{2n} \otimes C^2 + \beta^2 C^2 \otimes I_{2n}) + \beta(C \otimes C + C \otimes C) \\ &= (1 + \beta^2)(a^2 - b^2)I_{4n^2} + 2\beta C \otimes C = \phi I_{4n^2} + 2\beta C \otimes C. \end{aligned}$$

Using this equivalence for \mathcal{D}^2 , the formula for \mathcal{D}^{-1} is verified via

$$\begin{aligned} \mathcal{D}^{-1}\mathcal{D} &= \omega(2\phi I_{4n^2} - \mathcal{D}^2)\mathcal{D}^2 = \omega(2\phi I_{4n^2} - (\phi I_{4n^2} + 2\beta C \otimes C))(\phi I_{4n^2} + 2\beta C \otimes C) \\ &= \omega(\phi I_{4n^2} - 2\beta C \otimes C)(\phi I_{4n^2} + 2\beta C \otimes C) \\ &= \omega(\phi^2 I_{4n^2} - 4\beta^2 C^2 \otimes C^2) \\ &= \omega(\phi^2 - 4\beta^2(a^2 - b^2)^2)I_{4n^2}, \end{aligned}$$

which, by noting $\phi^2 - 4\beta^2(a^2 - b^2)^2 = \omega^{-1}$, is equivalent to I_{4n^2} . \square

LEMMA B.4. *Matrix \mathcal{Q}_2 from (7.20) is singular.*

Proof. Suppose that $\mathcal{Q}_2 w = 0$ if and only if $w = 0$, and so

$$\begin{bmatrix} 0 & \gamma^I \\ 0 & I \end{bmatrix} \otimes \begin{bmatrix} 0 & \gamma^I \\ 0 & I \end{bmatrix} w - \begin{bmatrix} I & 0 \\ \gamma^I & 0 \end{bmatrix} \otimes \begin{bmatrix} I & 0 \\ \gamma^I & 0 \end{bmatrix} w = 0 \Leftrightarrow \begin{bmatrix} 0 & \gamma^I \\ 0 & I \end{bmatrix} W \begin{bmatrix} 0 & 0 \\ \gamma^I & I \end{bmatrix} - \begin{bmatrix} I & 0 \\ \gamma^I & 0 \end{bmatrix} W \begin{bmatrix} I & \gamma^I \\ 0 & 0 \end{bmatrix} = 0,$$

where the equivalence holds by unvectorizing the first equation and $w = \text{vec}(W)$. This generalized Sylvester equation has a unique solution if and only if

$$\begin{bmatrix} 0 & \gamma^I \\ 0 & I \end{bmatrix} - \lambda \begin{bmatrix} I & 0 \\ \gamma^I & 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} I & \gamma^I \\ 0 & 0 \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 \\ \gamma^I & I \end{bmatrix}$$

are both regular matrix pencils and have no eigenvalues in common. However, as $\begin{bmatrix} 0 & \gamma^I \\ 0 & I \end{bmatrix}$ and $\begin{bmatrix} I & \gamma^I \\ 0 & 0 \end{bmatrix}$ are both singular, the two pencils share zero as an eigenvalue. \square

Acknowledgments. Many thanks to Mark Embree for the suggestion to look at the problem of computing Kreiss constants and to Michael L. Overton for pointing out the possible connection to computing the distance to uncontrollability and hosting the author at the Courant Institute in New York for several visits. The author is also grateful for the many useful suggestions from the referees, one of whom deserves special mention for pointing out both a more concise proof of Lemma 4.5 and the simpler explicit inverse formulas for \mathcal{B}_2 and $\tilde{\mathcal{B}}_2$ respectively given by (6.9) and (A.6).

REFERENCES

- [1] Z. BAI, J. DEMMEL, J. DONGARRA, A. RUHE, AND H. VAN DER VORST, EDS., *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*, Software Environ. Tools 11, SIAM, Philadelphia, 2000, <https://doi.org/10.1137/1.9780898719581>.
- [2] P. BENNER, V. MEHRMANN, AND H. XU, *A numerically stable, structure preserving method for computing the eigenvalues of real Hamiltonian or symplectic pencils*, Numer. Math., 78 (1998), pp. 329–358.
- [3] P. BENNER, V. MEHRMANN, AND H. XU, *A note on the numerical solution of complex Hamiltonian and skew-Hamiltonian eigenvalue problems*, Electron. Trans. Numer. Anal., 8 (1999), pp. 115–126.
- [4] P. BENNER AND T. MITCHELL, *Extended and improved criss-cross algorithms for computing the spectral value set abscissa and radius*, SIAM J. Matrix Anal. Appl., 40 (2019), pp. 1325–1352, <https://doi.org/10.1137/19M1246213>.
- [5] P. BENNER AND S. W. R. WERNER, *MORLAB – Model Order Reduction LABORatory (version 5.0)*, 2019, <http://www.mpi-magdeburg.mpg.de/projects/morlab>, <https://doi.org/10.5281/zenodo.3332716>.
- [6] P. BENNER AND S. W. R. WERNER, *Hankel-norm approximation of large-scale descriptor systems*, Adv. Comput. Math., 46 (2020), 40, <https://doi.org/10.1007/s10444-020-09750-w>.
- [7] R. B. BURCKEL, *An Introduction to Classical Complex Analysis. Vol. 1*, Pure Appl. Math. 64, Birkhäuser, Basel, 1979, <https://doi.org/10.1007/978-3-0348-9374-9>.
- [8] J. V. BURKE, A. S. LEWIS, AND M. L. OVERTON, *Robust stability and a criss-cross algorithm for pseudospectra*, IMA J. Numer. Anal., 23 (2003), pp. 359–375, <https://doi.org/10.1093/imanum/23.3.359>.
- [9] J. V. BURKE, A. S. LEWIS, AND M. L. OVERTON, *Pseudospectral components and the distance to uncontrollability*, SIAM J. Matrix Anal. Appl., 26 (2004), pp. 350–361, <https://doi.org/10.1137/S0895479803433313>.
- [10] E. K.-W. CHU, *The solution of the matrix equations $AXB - CXD = E$ and $(YA - DZ, YC - BZ) = (E, F)$* , Linear Algebra Appl., 93 (1987), pp. 93–105, [https://doi.org/10.1016/S0024-3795\(87\)90314-4](https://doi.org/10.1016/S0024-3795(87)90314-4).
- [11] T. A. DRISCOLL, N. HALE, AND L. N. TREFETHEN, *Chebfun Guide*, Pafnuty Publications, Oxford, UK, 2014, <http://www.chebfun.org/docs/guide/>.
- [12] R. EISING, *Between controllable and uncontrollable*, Systems Control Lett., 4 (1984), pp. 263–264, [https://doi.org/10.1016/S0167-6911\(84\)80035-3](https://doi.org/10.1016/S0167-6911(84)80035-3).
- [13] M. EMBREE AND B. KEELER, *Pseudospectra of matrix pencils for transient analysis of differential-algebraic equations*, SIAM J. Matrix Anal. Appl., 38 (2017), pp. 1028–1054, <https://doi.org/10.1137/15M1055012>.
- [14] M. GU, *New methods for estimating the distance to uncontrollability*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 989–1003, <https://doi.org/10.1137/S0895479897328856>.
- [15] M. GU, E. MENGI, M. L. OVERTON, J. XIA, AND J. ZHU, *Fast methods for estimating the distance to uncontrollability*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 477–502, <https://doi.org/10.1137/05063060X>.
- [16] N. GUGLIELMI AND M. L. OVERTON, *Fast algorithms for the approximation of the pseudospectral abscissa and pseudospectral radius of a matrix*, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 1166–1192, <https://doi.org/10.1137/100817048>.
- [17] B. KÄGSTRÖM AND L. WESTIN, *Generalized Schur methods with condition estimators for solving the generalized Sylvester equation*, IEEE Trans. Automat. Control, 34 (1989), pp. 745–751, <https://doi.org/10.1109/9.29404>.
- [18] M. KÖHLER, *Approximate Solution of Non-Symmetric Generalized Eigenvalue Problems and Linear Matrix Equations on HPC Platforms*, Dissertation, Otto-von-Guericke-Universität, Magdeburg, Germany, 2021, in preparation.
- [19] H.-O. KREISS, *Über die Stabilitätsdefinition für Differenzgleichungen die partielle Differentialgleichungen approximieren*, BIT, 2 (1962), pp. 153–181, <https://doi.org/10.1007/bf01957330>.
- [20] D. KRESSNER AND B. VANDEREYCKEN, *Subspace methods for computing the pseudospectral abscissa and the stability radius*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 292–313, <https://doi.org/10.1137/120869432>.
- [21] P. LANCASTER, *On eigenvalues of matrices dependent on a parameter*, Numer. Math., 6 (1964), pp. 377–387, <https://doi.org/10.1007/BF01386087>.
- [22] R. M. LARSEN, *PROPACK: Software for Large and Sparse SVD Calculations*, <http://sun.stanford.edu/~rmunk/PROPACK>.
- [23] A. J. LAUB, *Matrix Analysis for Scientists & Engineers*, SIAM, Philadelphia, 2005.

- [24] R. J. LEVEQUE AND L. N. TREFETHEN, *On the resolvent condition in the Kreiss matrix theorem*, BIT, 24 (1984), pp. 584–591, <https://doi.org/10.1007/BF01934916>.
- [25] E. MENGI, *Measures for Robust Stability and Controllability*, Ph.D. thesis, New York University, New York, NY 10003, 2006, https://cs.nyu.edu/media/publications/mengi_emre.pdf.
- [26] T. MITCHELL, *ROSTAPACK: RObust STAbility PACKage*, <http://timmitchell.com/software/ROSTAPACK>.
- [27] T. MITCHELL, *Fast Interpolation-Based Globality Certificates for Computing Kreiss Constants and the Distance to Uncontrollability*, preprint, <https://arxiv.org/abs/1910.01069>, 2019.
- [28] M. L. OVERTON AND R. S. WOMERSLEY, *Second derivatives for optimizing eigenvalues of symmetric matrices*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 697–718, <https://doi.org/10.1137/S089547989324598X>.
- [29] M. N. SPIJKER, *On a conjecture by LeVeque and Trefethen related to the Kreiss matrix theorem*, BIT, 31 (1991), pp. 551–555, <https://doi.org/10.1007/BF01933268>.
- [30] L. N. TREFETHEN AND M. EMBREE, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*, Princeton University Press, Princeton, NJ, 2005.
- [31] T. G. WRIGHT, *EigTool*, <http://www.comlab.ox.ac.uk/pseudospectra/eigtool/>, 2002.