

DUAL CERTIFICATES OF PRIMAL CONE MEMBERSHIP*

JOONYEOB LEE[†], DÁVID PAPP[‡], AND ANITA VARGA[§]

Abstract. We discuss easily verifiable cone membership certificates, that is, certificates proving relations of the form $b \in \mathcal{K}$ for convex cones \mathcal{K} , that consist of vectors in the dual cone \mathcal{K}^* . In the standard theory of duality, vectors in \mathcal{K}^* are associated with separating hyperplanes and interpreted as certificates of *non-membership* in \mathcal{K} . Complementing this, we present constructive certification schemes through which members of the dual cone can be interpreted as primal *membership* certificates. Every vector in the interior of \mathcal{K} is assigned a full-dimensional cone of certificates, making the numerical computation of rigorous certificates easy, provided that the dual cone has an efficiently computable logarithmically homogeneous self-concordant barrier. Of particular interest are cones that are low-dimensional linear images of much higher dimensional cones. In the context of optimization (as opposed to feasibility) problems, these certificates can be used to easily compute, using a closed-form formula, exact primal feasible solutions from suitable dual feasible solutions, with the guarantee that the closer the dual solutions are to optimality, the closer to optimality are the computed primal solutions, too. We demonstrate that the new certification scheme is applicable to virtually every tractable subcone of nonnegative polynomials commonly used in polynomial optimization (such as SOS, SONC, SAGE and SDSOS polynomials, among others), facilitating the computation of rigorous nonnegativity certificates using numerical algorithms.

Key words. Cone membership certificates, Self-concordant barrier functions, Nonsymmetric conic optimization, Duality, Interior-point algorithms, Polynomial optimization

MSC codes. 90C25, 90C51, 90C23, 49M29, 90C22

1. Introduction. Consider a closed, convex, full-dimensional, pointed cone $\mathcal{K} \subset \mathbb{R}^m$ and a vector $b \in \mathbb{R}^m$. If $b \notin \mathcal{K}$, then there is a hyperplane strongly separating b from \mathcal{K} : there exists a vector $y \in \mathbb{R}^m$ (the normal vector of this hyperplane) satisfying

$$y^T b < 0 < y^T z$$

for every $z \in \mathcal{K} \setminus \{0\}$ [13, Chapter 6]. This vector y can be interpreted as a *non-membership certificate*, that is, a proof that b does not belong to \mathcal{K} . By definition, y is an element of the interior of the dual cone of \mathcal{K} ,

$$\mathcal{K}^* = \{y \in \mathbb{R}^m \mid \forall x \in \mathcal{K} : x^T y \geq 0\};$$

this gives rise to the standard duality theory of convex (conic) programming. In this paper we show that with the help of self-concordance theory, *membership* (rather than non-membership) in the primal cone can be similarly certified by simply exhibiting suitable vectors from the dual cone. These certificates are not unique. On the contrary, elements of the interior of \mathcal{K} have a full-dimensional cone of membership certificates; this facilitates their computation using numerical methods and has important theoretical consequences.

Our first main contribution is that whenever \mathcal{K}^* is a hyperbolicity cone, membership in \mathcal{K} can be certified in a manner similar to the dual sums-of-squares certificates of [8] (we

***Funding:** This material is based upon work supported by the National Science Foundation under Grant No. DMS-1847865. This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-23-1-0370.

[†]North Carolina State University, Department of Mathematics. Email: jlee223@ncsu.edu.

[‡]North Carolina State University, Department of Mathematics. (ORCID: 0000-0003-4498-6417) Email: dpapp@ncsu.edu.

[§]North Carolina State University, Department of Mathematics. (ORCID: 0000-0001-7138-2921) Email: avarga@ncsu.edu.

call these *H-certificates* in [Definition 2.3](#)), but the generalization does not go any further ([Example 2.2](#)). Then we present a new certification scheme (termed *B-certificates*) that is equally efficient for the same cones as H-certificates, but can also be applied to all cones \mathcal{K} as long as \mathcal{K}^* is the domain of an efficiently computable logarithmically homogeneous self-concordant barrier function ([Theorem 2.5](#)).

We also present a unified algorithmic framework to efficiently compute these dual membership certificates. In [Sections 4 and 5](#), we argue that *short-step* interior-point algorithms relying on small neighborhoods are the natural approach to find such certificates.

Of particular interest are cones \mathcal{K}_A that are the linear images of a “simpler” but much higher dimensional cone \mathcal{K} . Suppose that

$$(1.1) \quad \mathcal{K}_A \stackrel{\text{def}}{=} \{Ax \mid x \in \mathcal{K}\},$$

where $A \in \mathbb{R}^{m \times n}$ with $m \ll n$ is a matrix with linearly independent rows. On the surface, solving an optimization problem over \mathcal{K}_A , or even checking whether a given vector b belongs to \mathcal{K}_A , requires the solution of an optimization problem over \mathcal{K} , and in order to certify $b \in \mathcal{K}_A$, we must present an $x \in \mathcal{K}$ satisfying $b = Ax$. This raises two issues: first, this “primal” membership certificate x is n -dimensional, potentially much larger than the vector whose membership is being certified, which may be prohibitive even when $m = \dim(\mathcal{K}_A)$ is small. Second, if a numerical method is used to compute x , the computed certificate is likely only an approximate certificate: $\|b - Ax\|$ is small but nonzero. If we need an exact certificate (e.g., a rational vector $x \in \mathcal{K} \cap \mathbb{Q}^n$ with $b = Ax$), additional post-processing is required to “round” or project the inexact numerical certificate to an exact one, taking care that this additional step does not result in the certificate leaving the cone \mathcal{K} , which is a highly non-trivial task in general, especially if A is ill-conditioned or the numerically computed x is close to the boundary of \mathcal{K} . For example, in the context of semidefinite programming representations of sums-of-squares cones, this problem was studied in detail in [\[31, 16\]](#).

In [Section 3](#) we demonstrate that our dual membership certificates overcome both of these issues. First, because the cone of certificates for every vector is full-dimensional, numerically computed certificates are automatically exact, rigorous certificates; and since they are just elements of \mathcal{K}_A^* , they are not higher dimensional than the certified vector. Second, the certification schemes come with a *reconstruction formula* for cones of the form [\(1.1\)](#) that allows the efficient computation of a primal membership certificate $x \in \mathcal{K}$ that by definition satisfies $b = Ax$ exactly. ([Theorem 3.2.](#)) For most cones \mathcal{K} that commonly arise in conic optimization problems, the formula automatically yields rational primal certificates x (as long as A is a rational matrix), which can also be verified efficiently in rational arithmetic.

In [Section 6](#), we show how dual certificates can be used to compute, with a closed-form formula, exact primal feasible solutions of conic optimization problems in standard form from suitable dual feasible solutions, with a guarantee that the closer the dual solutions are to optimality, the closer the computed primal feasible solutions will be to optimality. With this formula, most short-step, path-following IPAs (including dual-only methods) can be “retrofitted” to compute primal-dual feasible pairs with a guaranteed low duality gap with minimal additional computational cost in each iteration.

In [Section 7](#), we outline the application of these ideas in polynomial optimization. The most commonly used inner approximations of the cones of nonnegative polynomials (including sums-of-squares (SOS) polynomials, sums of nonnegative circuit (SONC) polynomials, sums of AM-GM exponentials (aka. SAGE functions), SDSOS polynomials and others) are all cones of the form [\(1.1\)](#) with an easily computable matrix $A \in \mathbb{Q}^{m \times n}$, usually with $m \ll n$, and a (possibly non-symmetric) cone \mathcal{K} for which \mathcal{K}^* admits a well-known logarithmically

homogeneous self-concordant barrier function. Therefore, the problem of rigorously certifying the nonnegativity of a polynomial using either of these cones amounts to the exact solution of a feasibility problem in \mathcal{K}_A , as discussed above. When \mathcal{K} is any of the aforementioned cones of polynomials, the dual certificates introduced in this paper can be interpreted as easily verifiable rigorous polynomial nonnegativity certificates that can be computed numerically without constructing an explicit “primal” representation (such as an explicit representation of an SOS polynomial as a sum of squared polynomials). Should such a representation be necessary, the reconstruction formula can be used to compute an exact representation in closed form using a numerically computed dual certificate.

Relation to prior work. The results in this paper are broad generalizations of prior work of Davis and Papp [8] on *weighted sums-of-squares* (WSOS) cones, which are convex cones instrumental in global polynomial optimization. Membership in these cones can be cast as a (feasibility) semidefinite program (SDP) [30]; every feasible solution to this SDP is an obvious membership certificate. However, the number of variables in this SDP is at least an order of magnitude larger than the dimension of the cone (and in the presence of a large number of weights, the ratio of the dimensions of these two cones can be arbitrarily large). Instead, the certificates introduced in [8] (simply termed there as *dual certificates*) are elements of the dual WSOS cone, and are thus low-dimensional, and can be more efficiently computed using non-symmetric cone programming methods than solving the conventional SDP. The theory presented in the present paper is a generalization of these earlier results, and the algorithms discussed here are simultaneously more general and more efficient (owing to their lower iteration complexity) than the algorithms presented in the earlier work.

In particular, the new certification scheme is applicable to virtually every tractable subcone of nonnegative polynomials commonly used in polynomial optimization besides WSOS polynomials; we outline some of these applications in Section 7.

Background: convex cones and self-concordant barriers. Central to the paper are the notions of *self-concordant barrier functions* and *logarithmic homogeneity*. The reader may find a comprehensive treatment of this subject in the monographs [25, 24, 34]; for easier referencing of key results, we recall a few fundamental definitions and properties below.

DEFINITION 1.1. Let $\mathcal{C} \subseteq \mathbb{R}^n$ be a convex cone with a non-empty interior \mathcal{C}° . A function $f : \mathcal{C}^\circ \rightarrow \mathbb{R}$ is called a strongly nondegenerate, self-concordant barrier (SCB) for \mathcal{C} if f is strictly convex, at least three times differentiable, and has the following three properties (with $\text{bd}(\mathcal{C})$ denoting the boundary of \mathcal{C}):

1. $f(x) \rightarrow \infty$ as $x \rightarrow \text{bd}(\mathcal{C})$ and
2. $|D^3 f(x)[h, h, h]| \leq 2 (D^2 f(x)[h, h])^{3/2}$.
3. Denoting the gradient and Hessian of f by g and H , respectively, we have $\nu \stackrel{\text{def}}{=} \sup_{x \in \mathcal{C}^\circ} g(x)^T H(x)^{-1} g(x) < \infty$.

We refer to the parameter ν above as the barrier parameter of f .

We say that f is a logarithmically homogeneous self concordant barrier (LHSCB) with barrier parameter ν , or a ν -LHSCB for short if, in addition to the properties above, it is also true that for every $x \in \mathcal{C}^\circ$ and $t > 0$, $f(tx) = f(x) - \nu \ln(t)$.

Given a positive definite, real, symmetric matrix M of order n , we denote by $\langle \cdot, \cdot \rangle_M$ the inner product $(x, y) \mapsto x^T M y$; the induced norm is $\|\cdot\|_M$. If $M = H(x)$ is the Hessian of an LHSCB at $x \in \mathcal{C}^\circ$, we may use the shorthands $\|\cdot\|_x$ in place of $\|\cdot\|_{H(x)}$ and $\|\cdot\|_x^*$ in place of $\|\cdot\|_{H(x)^{-1}}$. These are commonly referred to as the primal and dual *local norms* associated with the barrier function.

Proposition 1.2 below summarizes a few properties of self-concordant barrier functions used throughout the paper. The proofs of these statements can be found, for example, in [34, Chapters 2–3] and [24, Chapter 5].

PROPOSITION 1.2. *Suppose $f : \mathcal{C}^\circ \rightarrow \mathbb{R}$ is an SCB for some convex cone $\mathcal{C} \subseteq \mathbb{R}^n$, and let $g(x)$ and $H(x)$ denote its gradient and Hessian at $x \in \mathcal{C}^\circ$. Then the following hold:*

1. *The gradient map $g : \mathcal{C}^\circ \rightarrow \mathbb{R}^n$ is a bijection between \mathcal{C}° and $-(\mathcal{C}^*)^\circ$. That is, for every $b \in (\mathcal{C}^*)^\circ$ there exists a unique $y_b \in \mathcal{C}^\circ$ that satisfies $-g(y_b) = b$.*
2. *For every $x, y \in \mathcal{C}^\circ$,*

$$(1.2) \quad -g(x)^\top y \geq \|y\|_x.$$

If, in addition, f is a ν -LHSCB, then the following also hold:

3. *The gradient and Hessian satisfy*

$$g(tx) = t^{-1}g(x) \text{ and } H(tx) = t^{-2}H(x) \text{ for every } t > 0,$$

and

$$(1.3) \quad H(x)x = -g(x) \quad \text{and} \quad \|g(x)\|_x^* = \|x\|_x = \sqrt{-g(x)^\top x} = \sqrt{\nu}.$$

4. *If $A : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is an injective linear operator and $\mathcal{A} \stackrel{\text{def}}{=} \{x \in \mathbb{R}^m \mid A(x) \in \mathcal{C}^\circ\}$ is non-empty, then $x \mapsto f(A(x))$ is a ν -LHSCB for \mathcal{A} .*

2. Membership certificates from the dual cone. Throughout this section, we shall use the following assumptions and notational shortcuts: we assume that \mathcal{K} is a proper convex cone whose dual cone \mathcal{K}^* admits a ν -LHSCB f , and we denote the gradient and Hessian of f at $y \in (\mathcal{K}^*)^\circ$ by $g(y)$ and $H(y)$, respectively. Finally, we define

$$B(y) \stackrel{\text{def}}{=} H(y) + g(y)g(y)^\top.$$

The membership certification schemes introduced in this section are based on the following facts.

THEOREM 2.1. *Let \mathcal{K} be a proper cone whose dual cone \mathcal{K}^* is the domain of an SCB, and let g and H denote the gradient and Hessian of this barrier, respectively. Then for every $y \in (\mathcal{K}^*)^\circ$ and $z \in \mathcal{K}^*$,*

$$(2.1) \quad B(y)z \in \mathcal{K}.$$

If \mathcal{K}^ is the hyperbolicity cone of the hyperbolic polynomial p and f is the logarithmic barrier $-\ln(p(\cdot))$, then we also have*

$$(2.2) \quad H(y)z \in \mathcal{K}.$$

Proof. Using properties of the SCBs and the Cauchy–Schwarz inequality, we have that every $y \in (\mathcal{K}^*)^\circ$, $z \in \mathcal{K}^*$, and $w \in \mathcal{K}^*$ satisfy the following chain of inequalities:

$$(g(y)^\top w) (g(y)^\top z) \stackrel{(1.2)}{\geq} \|w\|_y \|z\|_y \geq |\langle w, z \rangle_y| = |w^\top H(y)z| \geq -w^\top H(y)z.$$

Rearranging the left- and right-hand sides, we get

$$0 \leq w^T H(y)z + (g(y)^T w) (g(y)^T z) = w^T (H(y) + g(y)g(y)^T) z = w^T B(y)z.$$

Since $w \in \mathcal{K}^*$ can be chosen arbitrarily, this yields (2.1), proving our first assertion.

The second claim is Güler's result from [12, Thm. 6.2]. \square

The assumption of hyperbolicity in the second half of Theorem 2.1 is necessary; the relation (2.2) does not hold for every cone.

EXAMPLE 2.2 (Relation (2.2) does not hold for every cone). *Consider the exponential cone \mathcal{E} defined by*

$$\mathcal{E} \stackrel{\text{def}}{=} \text{cl} \left(\left\{ (x_1, x_2, x_3) \in \mathbb{R}_+^2 \times \mathbb{R} \mid x_2 > 0, x_1 \geq x_2 e^{x_3/x_2} \right\} \right),$$

where $\text{cl}(\cdot)$ denotes (topological) closure. A well-known 3-LHSCB for \mathcal{E} is

$$f_{\mathcal{E}}(x_1, x_2, x_3) = -\ln x_1 - \ln x_2 - \ln(x_2 \ln(x_1/x_2) - x_3);$$

see, e.g., [23, Eq. (4.5)]. Now, it can be verified by direct computation that the following vectors all belong to \mathcal{E}° :

$$y = (6, 2, -3), \quad z = (2, 4, -3), \quad w = (416, 1, 6).$$

With some more tedious arithmetic, we can also compute that

$$w^T H(y)z = \frac{\ln(3)(3211 + 904 \ln(3)) - 4637}{9(3 + \ln(9))^2} < -0.075,$$

whose negativity proves that, although $y \in \mathcal{E}$ and $z \in \mathcal{E}$, $H(y)z \notin \mathcal{E}^*$. Therefore, $\mathcal{K} = \mathcal{E}^*$ does not satisfy the relation (2.2).

For another counterexample, consider the power cone $\mathcal{P}_{(2/3, 1/3)}$ defined by

$$\mathcal{P}_{(2/3, 1/3)} \stackrel{\text{def}}{=} \left\{ (x_1, x_2, x_3) \in \mathbb{R}_+^2 \times \mathbb{R} \mid (x_1^2 x_2)^{1/3} \geq |x_3| \right\}.$$

A well-known 3-LHSCB for this cone is

$$f_{\mathcal{P}_{(2/3, 1/3)}}(x_1, x_2, x_3) = -\ln \left((x_1^2 x_2)^{2/3} - x_3^2 \right) - \frac{1}{3} \ln(x_1) - \frac{2}{3} \ln(x_2);$$

see, e.g., [35]. Now, it can be verified by direct computation that the following vectors all belong to $\mathcal{P}_{(2/3, 1/3)}^\circ$:

$$y = (10, 1, 1), \quad z = (1, 20, 2), \quad w = (355, 1, 50),$$

however,

$$w^T H(y)z = \frac{12\,871 - 443\,620 \sqrt[3]{10} + 195\,500 \sqrt[3]{100}}{60 \left(1 - 10 \sqrt[3]{10}\right)^2} < -1.399,$$

proving that (2.2) is not satisfied.

Also note that the matrices $B(y)$ and $H(y)$ are automatically invertible for every $y \in (\mathcal{K}^*)^\circ$, since $H(y)$ is positive definite. We are now ready for our key definitions.

DEFINITION 2.3 (dual membership certificates). *Let $b \in \text{span}(\mathcal{K})$ be given. A vector $y \in (\mathcal{K}^*)^\circ$ is a B-certificate of b if, using the notation of [Theorem 2.1](#), $B(y)^{-1}b \in \mathcal{K}^*$. It is an H-certificate of b if $H(y)^{-1}b \in \mathcal{K}^*$. If the cone and the implied type of certificate are clear from the context, we may refer to these as dual certificates.*

[Theorem 2.1](#) justifies the terminology: if a vector $b \in \text{span}(\mathcal{K})$ has a B-certificate, then it is guaranteed to be in \mathcal{K} , and the same holds for H-certificates if \mathcal{K}^* is a hyperbolicity cone.

In both cases, the certificate y is easy to verify as long as the first two derivatives of the barrier function are easily computable (which in turn implies that membership in \mathcal{K}^* is easily verifiable): one needs to first verify $y \in \mathcal{K}^*$ and compute the matrix B or H at y ; the rest of the verification process amounts to simple linear algebra and verification of the membership of another vector in \mathcal{K}^* . In particular, many known LHSCBs have efficiently computable gradients and Hessians that are also rational at every rational point. In this case, the computed rational dual membership certificates can even be rigorously verified in rational arithmetic.

EXAMPLE 2.4. *Sums-of-squares (and weighted sums-of-squares) cones are subcones of polynomials that are non-negative over basic, closed semi-algebraic sets; they are well-known for their instrumental role in polynomial optimization [20]. Their dual cones are spectrahedral and thus hyperbolic, and this means that H-certificates are valid membership certificates for sums-of-squares polynomials. In fact, H-certificates generalize the concept of “dual certificates” introduced in the recent work [8] for cones of (weighted) sums-of-squares polynomials.*

When polynomials are represented in many of the commonly used bases (e.g., the monomial basis or the Chebyshev basis of the first or second kind), the duals of the corresponding sums-of-squares cones admit an LHSCB of the form $y \rightarrow -\ln \det \Lambda(y)$ with an injective linear map Λ that maps rational vectors to rational symmetric matrices. As a result, $g(y)$ is a rational vector and $H(y)$ is a rational matrix for every rational vector y . This yields a simple, independent proof of a theorem of Powers [32] on the existence of rational SOS certificates [8], and also means that rational H-certificates of SOS polynomials with rational coefficients can be efficiently verified in rational arithmetic [9]. We shall expound on this example in [Section 7.1](#).

We now show that the converse also holds: every vector in \mathcal{K}° has a B-certificate, and if the employed SCB is an LHSCB, then \mathcal{K}^* contains a *full-dimensional* subcone of vectors that certify a given $b \in \mathcal{K}^\circ$. The fact that these cones of certificates are full-dimensional is critical, as they allow us to compute them easily with numerical methods such those discussed in [Section 5](#).

THEOREM 2.5. *Let $b \in \mathcal{K}^\circ$, and consider its set of B- and H-certificates with respect to an LHSCB of $(\mathcal{K}^*)^\circ$:*

$$(2.3) \quad \mathcal{B}(b) \stackrel{\text{def}}{=} \{y \in (\mathcal{K}^*)^\circ \mid B(y)^{-1}b \in \mathcal{K}^*\}$$

and

$$(2.4) \quad \mathcal{H}(b) \stackrel{\text{def}}{=} \{y \in (\mathcal{K}^*)^\circ \mid H(y)^{-1}b \in \mathcal{K}^*\}.$$

Additionally, let y_b be the unique vector in $(\mathcal{K}^)^\circ$ that satisfies $-g(y_b) = b$. (Recall [Proposition 1.2](#).) Then $\mathcal{B}(b)$ and $\mathcal{H}(b)$ are full-dimensional subcones of $(\mathcal{K}^*)^\circ$ containing y_b .*

Proof. From the definitions in (2.3) and (2.4) and Proposition 1.2 it immediately follows that both $\mathcal{B}(b)$ and $\mathcal{H}(b)$ are subsets of $(\mathcal{K}^*)^\circ$ and that they are cones; all we have left to show is that they both contain every vector y in some neighborhood of y_b .

Recall from Eq. (1.3) that $H(y)^{-1}g(y) = -y$ for every $y \in (\mathcal{K}^*)^\circ$. Thus, we have

$$H(y_b)^{-1}b = -H(y_b)^{-1}g(y_b) = y_b \in (\mathcal{K}^*)^\circ,$$

proving that $y_b \in \mathcal{H}(b)$. And because the function $y \mapsto H(y)^{-1}b$ is continuous in some neighborhood of y_b , it is immediate that $y \in \mathcal{H}(b)$ for every y in some neighborhood of y_b , proving the second claim.

The proof of the first claim is similar, but we need to use, in addition, the Sherman–Morrison formula and the second identity of Eq. (1.3) to compute

$$(2.5a) \quad B(y)^{-1} = (H(y) + g(y)g(y)^\top)^{-1} = H(y)^{-1} - \frac{H(y)^{-1}g(y)g(y)^\top H(y)^{-1}}{1 + g(y)^\top H(y)^{-1}g(y)}$$

$$(2.5b) \quad \stackrel{(1.3)}{=} H(y)^{-1} - \frac{1}{1 + \nu} yy^\top,$$

after which we can see that

$$\begin{aligned} B(y_b)^{-1}b &= -B(y_b)^{-1}g(y_b) = -H(y_b)^{-1}g(y_b) + \frac{y_b (y_b^\top g(y_b))}{1 + \nu} \\ &\stackrel{(1.3)}{=} \left(1 - \frac{\nu}{\nu + 1}\right) y_b \in (\mathcal{K}^*)^\circ. \end{aligned}$$

That is, $y_b \in \mathcal{B}(b)$, moreover, every y in a neighborhood of y_b is also contained in $\mathcal{B}(b)$. \square

We close this section by showing that a B-certificate is always an H-certificate, but not vice versa. That is, the cone of H-certificates for a given vector is generally larger than the cone of B-certificates (but, unlike B-certificates, they only certify membership when \mathcal{K}^* is the hyperbolicity cone).

LEMMA 2.6. *If a vector $y \in \mathcal{K}^*$ is a B-certificate for b , then y is also an H-certificate for b . The converse, in general, does not hold.*

Proof. Assume $B(y)^{-1}b \in \mathcal{K}^*$. From our formula (2.5) for $B(y)^{-1}$, we see that

$$H(y)^{-1}b = B(y)^{-1}b + \frac{y^\top b}{1 + \nu} y.$$

By assumption, the first term, $B(y)^{-1}b$, belongs to \mathcal{K}^* . So does the second term, $\frac{y^\top b}{1 + \nu} y$, since $y \in \mathcal{K}^*$ by assumption and $b \in \mathcal{K}$, therefore $y^\top b \geq 0$. We conclude that $H(y)^{-1}b \in \mathcal{K}^*$, that is, y is an H-certificate of b , as claimed.

Example 2.2 shows that the converse in general does not hold. For some cones that are not hyperbolicity cones, H-certificates exist even for some vectors $b \notin \mathcal{K}$ (that is, they are not membership certificates in general cones), while we have seen that B-certificates only exist for vectors in \mathcal{K} . \square

3. A reconstruction formula for linear images of cones. We now turn our attention to our primary motivating examples, which are surjective linear images of high-dimensional cones. Suppose that

$$(3.1) \quad \mathcal{K}_A = \{Ax \mid x \in \mathcal{K}\}, \text{ and therefore } \mathcal{K}_A^* = \{y \mid A^\top y \in \mathcal{K}^*\},$$

the rows of A are linearly independent, and assume that \mathcal{K}^* has a non-empty interior that is the domain of a known, easily computable LHSCB. Even if \mathcal{K} has a known LHSCB, its linear image \mathcal{K}_A does not “inherit” an obvious barrier function; this is the main reason why it is difficult to optimize over \mathcal{K}_A , and why deciding feasibility in \mathcal{K}_A is non-trivial even if deciding feasibility in \mathcal{K} is easy. However, because the A^T map is injective, if \mathcal{K}_A^* is full-dimensional, it inherits an easily computable LHSCB from \mathcal{K}^* : if f is an LHSCB for \mathcal{K}^* , then

$$(3.2) \quad f_A \stackrel{\text{def}}{=} y \mapsto f(A^T y)$$

is automatically an LHSCB with domain $(\mathcal{K}_A^*)^\circ$. (Recall Proposition 1.2.) The gradient and Hessian of f_A can be expressed in closed form.

LEMMA 3.1. *Let g and H denote the gradient and Hessian of f . Then the gradient g_A and Hessian H_A of the function f_A defined in (3.2) can be calculated as follows*

$$(3.3) \quad g_A(y) = Ag(A^T y) \quad H_A(y) = AH(A^T y)A^T.$$

In the same vein,

$$(3.4) \quad B_A(y) = AB(A^T y)A^T.$$

THEOREM 3.2 (Reconstruction Theorem). *Suppose $b \in \mathcal{K}_A^\circ$, and let $y \in \mathcal{B}(b)$ be one of its dual membership certificates. Then the vector*

$$(3.5) \quad x = B(A^T y)A^T (AB(A^T y)A^T)^{-1} b$$

satisfies $Ax = b$ and $x \in \mathcal{K}$.

The same statement holds for H -certificates if in (3.5) we replace every instance of the matrix B with H .

Proof. We detail the proof for B -certificates only; it can be translated to H -certificates word-for-word by, as suggested in the statement, substituting the Hessian H in place of B throughout the proof.

We first note that x is well-defined: B is well-defined and invertible, since $y \in (\mathcal{K}_A^*)^\circ$ and $A^T y \in (\mathcal{K}^*)^\circ$; additionally, A is of full row-rank, which ensures that the matrix $(AB(A^T y)A^T)$ inverted in the formula is indeed invertible.

The first of our two claims, $Ax = b$, is immediate from basic arithmetic.

To show that $x \in \mathcal{K}$, first note that we can apply Eqs. (3.3)–(3.4) and Theorem 2.1 to the cone \mathcal{K}_A^* and its LHSCB f_A to see that

$$B_A(y)^{-1}b = (AB(A^T y)A^T)^{-1}b \in \mathcal{K}_A^*.$$

Then, by definition, we have

$$A^T(AB(A^T y)A^T)^{-1}b \in \mathcal{K}^*,$$

and using Theorem 2.1 again, we obtain the desired relation

$$B(A^T y)A^T(AB(A^T y)A^T)^{-1}b \in \mathcal{K}.$$

□

This theorem gives a closed-form formula with which we can compute a suitable pre-image of $b \in \mathcal{K}_A$ that automatically belongs to \mathcal{K} to make the verification of $b \in \mathcal{K}_A$ even simpler (as long as membership in \mathcal{K} can be easily verified). In other words, the formula (3.5) converts the dual certificate y into a conventional, “primal” certificate of b ’s membership in \mathcal{K}_A . We note, however, that the relation $b \in \mathcal{K}_A$ is already rigorously certified by the relation $B_A(y)^{-1}b \in \mathcal{K}_A^*$, and the computation of the primal certificate x is not necessary to produce or verify a dual certificate. On the contrary, dual certificates allow us to optimize directly over \mathcal{K}_A and verify the feasibility of the computed (near-)optimal solutions without resorting to computing and verifying the membership of any vector in \mathcal{K} , which is useful whenever $\dim(\mathcal{K}_A) \ll \dim(\mathcal{K})$. We detail the algorithmic aspects of these questions in Section 5.

Primal reconstruction and weighted least-squares. The formula (3.5) can be interpreted as the optimal solution of the equality-constrained (weighted) least-squares problem

$$(3.6) \quad \min_x \left\{ \left\| -g(A^T y) - x \right\|_{B(A^T y)^{-1}} \mid Ax = b \right\},$$

where $y \in (\mathcal{K}^*)^\circ$ is a given vector. Straightforward arithmetic involving (3.3)–(3.4) (without any reliance on Proposition 1.2) shows that if λ denotes the Lagrange multipliers associated with the equality constraints, then the optimality conditions of (3.6) yield the optimal multipliers

$$\lambda^* = B_A(y)^{-1} (-g_A(y) - b)$$

and the optimal (primal) solution

$$(3.7) \quad x^* = B(A^T y) A^T B_A(y)^{-1} b - (I - B(A^T y) A^T B_A(y)^{-1} A) g(A^T y).$$

Using (1.3), the second term simplifies to 0, after which the equation (3.7) reduces to (3.5).

The interpretation of (3.6) is fairly clear. The range of $-g$ is the primal cone \mathcal{K}° , in which we wish x to belong (recall the first claim of Proposition 1.2); thus, this problem seeks to find an x satisfying $Ax = b$ that is close (in a suitable, y -dependent norm) to a given vector that is already in \mathcal{K}° .

For further insight, note that the optimal value of (3.6) is

$$\left\| -g(A^T y) - x^* \right\|_{B(A^T y)^{-1}} = \left\| -g_A(y) - b \right\|_{B_A(y)^{-1}},$$

and analogously, for H-certificates we may replace every occurrence of B with H in the previous calculations to obtain the optimal objective value $\left\| -g_A(y) - b \right\|_y^*$. We can use these observations to paint a more geometric picture, and to derive easy-to-use sufficient conditions for dual certification. For this, we briefly recall the notion of Dikin ellipsoids.

DEFINITION 3.3. *Given a proper convex cone $\mathcal{C} \subseteq \mathbb{R}^n$ whose interior is the domain of an LHSCB, we define the local (open) ball centered at $x \in \mathcal{C}^\circ$ with radius r as $\mathcal{B}(x, r) \stackrel{\text{def}}{=} \{v \in \mathbb{R}^n \mid \|v - x\|_x < r\}$. These balls are called Dikin ellipsoids. In the same vein, for a $y \in (\mathcal{C}^*)^\circ$ we have the dual Dikin ellipsoid of radius r defined as $\mathcal{B}^*(y, r) \stackrel{\text{def}}{=} \{v \in \mathbb{R}^n \mid \|v - y\|_y^* < r\}$.*

It is well-known that Dikin ellipsoids of radius 1 stay entirely within their respective cones [34, Chapter 2]. Applying this to \mathcal{K}_A and \mathcal{K} , we see that if

$$(3.8) \quad \min_x \left\{ \left\| -g(A^T y) - x \right\|_{H(A^T y)^{-1}} \mid Ax = b \right\} = \left\| -g_A(y) - b \right\|_y^* < 1,$$

then $b \in \mathcal{K}_A$, moreover, if \mathcal{K}^* is a hyperbolicity cone, this is also proven by the fact that its preimage x given by the reconstruction formula belongs to \mathcal{K} . That is, the condition (3.8) is a sufficient (but not necessary!) condition for a dual vector to be an H-certificate of b .

We can state this result more generally, for arbitrary convex cones instead of cones of the form \mathcal{K}_A . In the next section, we explore this connection between Dikin ellipsoids and H- and B-certificates further.

4. Dikin ellipsoids and dual certificates. Throughout this section, we consider again a general proper convex cone \mathcal{K} , assuming only that $(\mathcal{K}^*)^\circ$ is the domain of a ν -LHSCB. As before, the gradient and Hessian of this barrier function are denoted by g and H , respectively.

THEOREM 4.1. *Suppose $y \in (\mathcal{K}^*)^\circ$ satisfies the relation $\| -g(y) - b \|_y^* < 1$ or (equivalently, using the primal local norm) $\| y - H(y)^{-1}b \|_y < 1$. Then y is an H-certificate for $b \in \mathcal{K}$.*

Proof. Identical to the last argument of the previous section. \square

The converse is not true in general, that is, the cone generated by Dikin ellipsoids may be a proper subcone of the cone of H-certificates. We can also use Dikin ellipsoids to derive analogous sufficient conditions for a vector to be a B-certificate:

THEOREM 4.2. *Let \mathcal{K}^* be the domain of a ν -LHSCB, and suppose that $b \in \mathcal{K}$ and $y \in (\mathcal{K}^*)^\circ$ satisfy the relation $b \in \mathcal{B}^*(-g(\delta y), r)$ for some $\delta > 0$ and $r = \frac{1}{2(\nu+1)}$. Then y is a B-certificate of b .*

Proof. It suffices to prove our theorem for $\delta = 1$, since $\mathcal{B}(b)$ is a cone. We will show our claim by arguing that in this case,

$$\lambda = \frac{1 + \nu}{1 + \nu - y^T b}$$

is a positive scalar for which $\| y - \lambda B(y)^{-1}b \|_y < 1$, therefore y is a B-certificate for λb .

To lighten the notation, we shall use the shorthands g , H , and B for $g(y)$, $H(y)$, and $B(y)$, respectively.

Note that

$$0 \leq \| y - H^{-1}b \|_y^2 = y^T H y + b^T H^{-1}b - 2y^T b.$$

By the definition of the dual local norm and the fact that $y^T H y \stackrel{(1.3)}{=} \nu$, we get

$$(4.1) \quad y^T b \leq \frac{1}{2}(y^T H y + b^T H^{-1}b) = \frac{1}{2}(\nu + \| b \|_y^{*2}).$$

Also by the assumption $b \in \mathcal{B}^*(-g(y), r)$ and the fact that $\| g \|_y^* = \sqrt{g^T H^{-1}g} = \sqrt{\nu}$ we obtain

$$\begin{aligned} \| b \|_y^* &\leq \| -g \|_y^* + \| b - (-g) \|_y^* \\ &= \| -g \|_y^* + \| y - H^{-1}b \|_y \\ &\leq \sqrt{\nu} + r. \end{aligned}$$

Therefore, (4.1) implies

$$y^T b \leq \frac{1}{2}(\nu + \| b \|_y^{*2}) \leq \frac{1}{2}(\nu + (\sqrt{\nu} + r)^2) = \nu + \sqrt{\nu}r + \frac{1}{2}r^2,$$

and we get

$$1 + \nu - y^T b \geq 1 - \sqrt{\nu}r - \frac{1}{2}r^2 = 1 - \frac{\sqrt{\nu}}{2(1+\nu)} - \frac{1}{8(1+\nu)^2} \geq \frac{23}{32} > 0.$$

Thus we can conclude that λ is positive; indeed

$$0 < \lambda \leq \frac{32}{23}(1+\nu).$$

From this, using (2.5) and the definition of λ , we also obtain

$$\begin{aligned} \|y - B^{-1}\lambda b\|_y &= \left\| y - H^{-1}\lambda b + \frac{\lambda}{1+\nu} y y^T b \right\|_y \\ &= \left\| \left(1 + \frac{\lambda}{1+\nu} y^T b \right) y - H^{-1}\lambda b \right\|_y \\ &= \lambda \|y - H^{-1}b\|_y \\ &\leq \frac{32}{23} \frac{1+\nu}{2(1+\nu)} < 1, \end{aligned}$$

completing the proof. \square

We underscore that [Theorem 4.2](#) is only a sufficient, but not necessary, condition for a vector to be a B-certificate.

5. Computing dual membership certificates. We now turn to the question of efficiently deciding whether $b \in \mathcal{K}^\circ$ holds and computing an H- or B-certificate if it does. In addition to assuming that $(\mathcal{K}^*)^\circ$ is the domain of an efficiently computable LHSCB and that we have a known $y \in (\mathcal{K}^*)^\circ$, we also assume that a vector $w \in \mathcal{K}^\circ$ is given. (The latter is essentially without loss of generality; e.g., we can always choose $w = -g(y) \in \mathcal{K}^\circ$.) Given such a w , the vector $b - \alpha w$ also belongs to \mathcal{K}° for every sufficiently negative α . Furthermore, $b \in \mathcal{K}$ holds if and only if there exists a nonnegative α such that $b - \alpha w \in \mathcal{K}$. Based on this observation, we can consider the following optimization problem and its dual:

$$(5.1) \quad \begin{array}{ll} \max_{z, \alpha} & \alpha \\ \text{s.t.} & \alpha w + z = b \\ & z \in \mathcal{K}, \alpha \in \mathbb{R} \end{array} \quad \begin{array}{ll} \min_y & b^T y \\ \text{s.t.} & w^T y = 1 \\ & y \in \mathcal{K}^* \end{array}$$

It is clear from our assumptions that strong duality holds with attainment on both the primal and dual side, since both problems have a Slater point. The optimal objective function value of (5.1) is positive if and only if $b \in \mathcal{K}^\circ$. However, standard duality theory and an approximately optimal numerical solution of *either problem* cannot rigorously certify this fact: a strictly feasible dual solution y can be scaled to exactly satisfy the only linear constraint $w^T y = 1$, and thus we can obtain a true, verifiable dual feasible solution, however, the corresponding objective value is an *upper bound* on the exact optimal value, therefore, a nonnegative dual value implies nothing. On the primal side we get a lower bound, and so an exact primal feasible (z, α) with $\alpha \geq 0$ would be a rigorous certificate, but numerical methods for problems of this form (IPAs in particular) tend to yield only approximately feasible primal solutions, meaning that only a vector close to b is certified to be in the cone.

In contrast, any $y \in \mathcal{B}(b)$ (or $y \in \mathcal{H}(b)$ if \mathcal{K}^* is a hyperbolicity cone) does rigorously certify $b \in \mathcal{K}$. Our goal is to compute such a dual certificate y .

The problem (5.1) is a *nonsymmetric convex conic optimization problem*, with the particular difficulty that no assumptions are made about the cone \mathcal{K} . There are only a handful of primal-dual interior-point algorithms (IPAs) developed for this general class of problems, which require only one LHSCB but simultaneously compute near-optimal primal and dual solutions; see, for example, [37, 29, 27, 6].¹ All of these are suitable tools for numerically addressing the decision problem of whether $b \in \mathcal{K}$. However, if we also aim to obtain dual certificates, it appears that short-step IPAs that operate within a close neighborhood of the central path are the most appropriate methods. This is due to the relationship between the definition of this neighborhood and the sufficient conditions we derived for dual certificates in Theorems 4.1 and 4.2.

To make this more precise, consider for a moment a general conic optimization problem in standard form

$$(5.2) \quad \begin{array}{ll} \inf_x & c^\top x \\ \text{s.t.} & Ax = b \\ & x \in \mathcal{K} \end{array} \quad \begin{array}{ll} \sup_y & b^\top y \\ \text{s.t.} & A^\top y + s = c \\ & s \in \mathcal{K}^* \end{array}$$

whose strict feasible set is

$$\mathcal{F}^\circ \stackrel{\text{def}}{=} \{(x, y, s) \in \mathcal{K}^\circ \times \mathbb{R}^m \times (\mathcal{K}^*)^\circ \mid Ax = b, A^\top y + s = c\}.$$

In our recent paper [27], we proposed IPAs for solving (5.2) that requires only an LHSCB for the *primal* cone \mathcal{K} and operates in the neighborhood

$$(5.3) \quad \mathcal{N}(\eta, \tau) = \{(x, y, s) \in \mathcal{F}^\circ \mid \|s + \tau g(x)\|_x^* \leq \eta\tau\},$$

where $\eta \in (0, 1)$ is the neighborhood radius (a fixed parameter throughout the algorithm, chosen by the user) and τ is the central path parameter (which tends geometrically to 0 as we approach optimality). Similar definitions appear, for example, in [24, 36, 37].

When specialized to (5.1) (keeping in mind that \mathcal{K}^* plays the role of $\mathcal{K}!$), the strict feasible region and neighborhood defined above simplify to

$$\begin{aligned} \mathcal{F}^\circ &= \{(z, \alpha, y) \in \mathcal{K}^\circ \times \mathbb{R} \times (\mathcal{K}^*)^\circ \mid \alpha w + z = b, w^\top y = 1\} \text{ and} \\ \mathcal{N}(\eta, \tau) &= \{(z, \alpha, y) \in \mathcal{F}^\circ \mid \|y - H(y)^{-1}(z/\tau)\|_y \leq \eta\}. \end{aligned}$$

Recall that in Section 4, we derived the following sufficient conditions for B- and H-certificates, which are nearly identical to the formula defining \mathcal{N} :

$$\begin{aligned} \|y - H(y)^{-1}(b/\tau)\|_y &< 1 && \implies y \in \mathcal{H}(b) \text{ and} \\ \|y - H(y)^{-1}(b/\tau)\|_y &< \frac{1}{2(1+\nu)} && \implies y \in \mathcal{B}(b) \end{aligned}$$

¹In these works it is assumed that the *primal* cone \mathcal{K} is the domain of a tractable LHSCB, with no assumptions on the *dual* cone \mathcal{K}^* . We made the opposite assumption for (5.1) to maintain consistency with the rest of this paper. Since these methods solve the primal and dual problems simultaneously, this is only a notational difference.

for any $\tau > 0$. Therefore, if we have a strictly feasible solution $(z, \alpha, y) \in \mathcal{F}^\circ$ to the optimization problem (5.1) with $(z, \alpha, y) \in \mathcal{N}(\eta, \tau)$ for any $\tau > 0$ and $\eta < 1$, then the dual feasible solution y is an H-certificate for $b - \alpha w$. If $\eta < 1/(2(1 + \nu))$ holds, then y is a B-certificate for $b - \alpha w$. Together with $\alpha \geq 0$, this certifies $b \in \mathcal{K}$ in both cases.

In conclusion, if $b \in \mathcal{K}^\circ$ and we use the IPAs proposed in [27] (or any other polynomial IPA operating in a similar Dikin-ellipsoid based neighborhood of the central path) to solve (5.1), we can efficiently compute a dual membership certificate.

Another desirable property of this approach is that even dual feasible solutions to (5.1) that are far from optimal (but are near the central path) can be meaningful, interpreting α as the distance from the boundary of the cone in the direction w , since every iterate (z, α, y) of a feasible IPA that operates within the neighborhood defined above is a rigorously certified bound α on the optimal value, with a dual certificate y . For instance, in the polynomial optimization problems considered in Section 7, we can choose the constant 1 polynomial as w and compute a sequence of rigorously certified lower bounds on a given polynomial. (We elaborate on this application in Section 7.)

6. Near-optimal primal feasible solutions. We now turn our attention from certifying membership (and the associated optimization problem (5.1)) to general convex optimization problems in standard conic form (5.2).

The main result of this section is that as long as we employ short-step dual or primal-dual IPAs to solve these problems (an even broader class of algorithms than those considered in the previous section), *an exact feasible solution x to the primal problem can be computed in closed form from any dual feasible solution y near the central path*, with a guarantee that *the closer y is to optimality, the closer the computed x is to optimality*.

To see how this can be algorithmically useful, besides being of independent theoretical interest, let us reformulate (5.2) using the notation

$$\mathcal{K}_{c,A} \stackrel{\text{def}}{=} \left\{ \begin{pmatrix} c^T x \\ -Ax \end{pmatrix} \mid x \in \mathcal{K} \right\}$$

and an auxiliary variable γ , as follows:

$$(6.1) \quad \begin{array}{ll} \inf_{\gamma \in \mathbb{R}} & \gamma \\ \text{s.t.} & \begin{pmatrix} \gamma \\ -b \end{pmatrix} \in \mathcal{K}_{c,A} \end{array} \quad \begin{array}{ll} \sup_{y \in \mathbb{R}^m} & b^T y \\ \text{s.t.} & \begin{pmatrix} 1 \\ y \end{pmatrix} \in \mathcal{K}_{c,A}^* \end{array}$$

The problems (5.2) and (6.1) are almost equivalent: they have the same optimal values, and their dual feasible regions are also clearly the same; however, (6.1) eliminates the variables x and s , at the cost that the new primal optimization problem is now over a cone $\mathcal{K}_{c,A}$, for which we may not have an efficiently computable LHSCB even if \mathcal{K} has one. The dual problem also involves a new cone, $\mathcal{K}_{c,A}^*$. If A has many more columns than rows, then the formulation (6.1) involves cones of much lower dimension, and may be solved more efficiently, than the original problem (5.2).

In addition to our previous assumption that $(\mathcal{K}^*)^\circ$ is the domain of a known ν -LHSCB $f_{\mathcal{K}^*}$, let us make the common regularity assumptions that (1) both the primal and dual problems are feasible (therefore we have finite optimal values); (2) $\begin{pmatrix} c^T \\ -A \end{pmatrix}$ has linearly independent rows; and (3) $\mathcal{K}_{c,A}^*$ has a nonempty interior. The last two assumptions ensure

that $(\mathcal{K}_{c,A}^*)^\circ$ is also the domain of a known ν -LHSCB (Proposition 1.2, Item 4.), namely

$$(6.2) \quad f_{\mathcal{K}_{c,A}^*}(y_0, y) \stackrel{\text{def}}{=} f_{\mathcal{K}^*}(cy_0 - A^T y).$$

Now, suppose that we have an algorithm that we can apply to (6.1) and compute, using only the dual barrier $f_{\mathcal{K}_{c,A}^*}$, the (approximately) optimal primal and dual *values* of the problems (5.2) and (6.1), and a (near-)optimal dual feasible solution y for these problems without computing any primal feasible solutions of (5.2). When A has many more columns than rows, circumventing the iterative optimization of x in this manner can save computational time and memory [28, 27], but many applications call for the actual primal solution to (5.2), not just the optimal value. As we shall see, we can still apply the same methods, and a primal feasible and (near-)optimal solution to the original problem (5.2) can be computed in post-processing using our reconstruction formula, at little additional cost. In fact, throughout the algorithms, primal feasible solutions that serve as easily verifiable certificates of our progress towards optimality can be computed in each iteration at a negligible additional cost.

Mirroring our discussion in the Introduction on feasibility certificates, one could argue that these reconstructed feasible solutions are preferable to the numerically computed ones, since they are theoretically guaranteed to be exact if computed (as post-processing) in exact arithmetic. Alternatively, one may use this throughout the algorithm as a guarantee that the numerical method did not veer too far from optimality due to numerical problems (or more reliably detect numerical problems), without resorting to a method that runs entirely in exact or high-precision arithmetic.

In the remainder of this section, we cover primal-dual and dual-only methods separately, as they rely on slightly different neighborhood definitions.

6.1. Exact primal reconstruction in primal-dual IPAs. In this section, we discuss how exact primal feasible solutions can be reconstructed from dual feasible solutions in the context of the authors' primal-dual path-following IPA in [27] and other methods that operate in the neighborhood (5.3).

The arguments for this case are essentially identical to the ones used in Section 5, after noting that (6.1) can also be equivalently written in the form of (5.1), with γ playing the role of α , the vector $e_0 = (1, 0, \dots, 0)$ playing the role of w , and $(0, -b)$ playing the role of A . (The only difference is that e_0 does not necessarily belong to $(\mathcal{K}_{c,A})^\circ$, but this is irrelevant now.)

With this change in notation, the summary of the previous section is that the application of IPAs using the neighborhood (5.3) to the optimization problem (6.1) yields a sequence of iterates (γ, y) , each of which has the property that $(1, y)$ certifies the relation $(\gamma, -b) \in \mathcal{K}_{c,A}$, that is, it certifies the existence of an $x \in \mathcal{K}$ satisfying $Ax = b$ and $c^T x = \gamma$ without ever computing such an x in the interim. We may compute such a primal vector x using our reconstruction theorem (Theorem 3.2), with $f_{\mathcal{K}_{c,A}^*}$ in place of f , $(1, y)$ in place of y and $(\gamma, -b)$ in place of b . And since these IPAs require the computation of the gradient and the factorization of the Hessian of the barrier in each iteration, the additional computation represented by the reconstruction formula (3.5) does not add a substantial computational cost. The bulk of the additional cost is due to matrix-vector multiplications that are cheaper than a single iteration of the IPA.

6.2. Exact primal reconstruction in dual IPAs. It may appear unsurprising that a primal feasible solution x of (5.2) can be reconstructed from a suitable corresponding primal

feasible solution of (6.1), but we emphasize that this is only made possible with a suitable dual certificate y . To further underscore this point, and to make our results more generally applicable, in this section we show that *dual-only* methods applied to (6.1) that make no reference to the primal cone $\mathcal{K}_{c,A}$ can also be “retrofitted” to compute exact primal feasible solutions from suitable dual feasible solutions. Our approach is applicable to all short-step, path-following, dual IPAs (or primal IPAs applied to our dual problem).

When applied to (6.1), such methods generate a sequence of dual vectors y in a neighborhood of the *dual central path* $\{y_\tau \mid \tau > 0\}$, where y_τ is the unique minimizer of the problem

$$(6.3) \quad \min \left\{ \tau f_{\mathcal{K}_{c,A}^*}(y_0, y) - b^T y \mid y_0 = 1, (y_0, y) \in (\mathcal{K}_{c,A}^*)^\circ \right\},$$

and the neighborhood is defined in terms the local norm associated with the dual barrier. Often, the central path parameter τ in (6.3) is equal to or bounded from below by the scaled duality gap $\mu = (c^T x - b^T y)/\nu$ at each iteration, and, crucially, $b^T y_\tau$ is within $\nu\tau$ of the optimal objective function value – this is what guarantees linear convergence of the methods in the objective function. See [24, Theorem 5.3.10] for a classic example. We have the following general scheme for primal reconstruction in dual methods that follow the above principle.

THEOREM 6.1. *Assume that $(1, y) \in (\mathcal{K}_{c,A}^*)^\circ$ (equivalently, that (y, s) is a dual feasible solution to (5.2) for some $s \in \mathcal{K}^\circ$), and let $\gamma \in \mathbb{R}$ be an arbitrary desired primal objective function value. Denote the gradient and Hessian of the ν -LHSCB $f_{\mathcal{K}^*}$ by $g_{\mathcal{K}^*}$ and $H_{\mathcal{K}^*}$, and let $B_{\mathcal{K}^*} = H_{\mathcal{K}^*} + g_{\mathcal{K}^*} g_{\mathcal{K}^*}^T$. Then the vector*

$$(6.4) \quad x_\gamma \stackrel{\text{def}}{=} B_{\mathcal{K}^*}(v) A_c^T (A_c B_{\mathcal{K}^*}(v) A_c^T)^{-1} b_\gamma,$$

where

$$A_c = \begin{pmatrix} c \\ -A \end{pmatrix}, \quad b_\gamma = \begin{pmatrix} \gamma \\ -b \end{pmatrix}, \quad \text{and} \quad v = A_c^T \begin{pmatrix} 1 \\ y \end{pmatrix}$$

satisfies $Ax_\gamma = b$ and $c^T x_\gamma = \gamma$. If, in addition, y is sufficiently close to the unique minimizer y_τ of (6.3) and we choose $\gamma = b^T y + \nu\tau$, then $x_\gamma \in \mathcal{K}$, that is, x_γ is a primal feasible solution of (5.2) whose objective value is no more than $\nu\tau$ away from optimal.

As in the rest of this paper, if \mathcal{K}^* is hyperbolic and $f_{\mathcal{K}^*}$ is the corresponding logarithmic barrier, then every instance of $B_{\mathcal{K}^*}$ in (6.4) can be replaced by $H_{\mathcal{K}^*}$.

Proof. The optimality conditions of (6.3) yield the following partial characterization of y_τ : the last m components of the the gradient $g(\cdot)$ of $f_{\mathcal{K}_{c,A}^*}$ are

$$\left. \frac{\partial f}{\partial y} f_{\mathcal{K}_{c,A}^*}(1, y) \right|_{y=y_\tau} = b/\tau.$$

To determine the missing zeroth component of the gradient, $g_0 = \frac{\partial f}{\partial y_0} f_{\mathcal{K}_{c,A}^*}$, we once again rely on Eq. (1.3) from Proposition 1.2. We have

$$\nu = - \begin{pmatrix} g_0 \\ b/\tau \end{pmatrix}^T \begin{pmatrix} 1 \\ y_\tau \end{pmatrix} = -g_0 - \tau^{-1} b^T y_\tau,$$

therefore, the barrier gradient satisfies

$$(6.5) \quad -\tau g(1, y_\tau) = \begin{pmatrix} b^T y_\tau + \tau\nu \\ -b \end{pmatrix}.$$

Recall from the proof of [Theorem 2.5](#) that [Eq. \(6.5\)](#) implies that $(1, y_\tau)$ is both an H- and B-certificate of the membership of the right-hand side vector in $\mathcal{K}_{c,A}^*$, certifying the existence of a feasible solution x in [\(5.2\)](#) whose attained value is $c^\top x = b^\top y_\tau + \tau\nu$. From the same theorem, every vector y sufficiently close to y_τ has the same property; that is, $(1, y)$ certifies the existence of a primal feasible solution x whose objective function value is $b^\top y + \tau\nu$.²

We may construct such a primal vector explicitly using our reconstruction theorem ([Theorem 3.2](#)), with A_c in place of A , v in place of y , and the vector b_γ in place of b . The resulting formula is precisely [Eq. \(6.4\)](#). \square

We note that because the reconstructed x_γ only depends on γ through the term b_γ , the vector x_γ is an affine function of γ that can be computed as easily as if γ was a known scalar. (See [Example 6.2](#) for an illustration.) Thus, we can apply [Theorem 6.1](#) “parametrically” to certify, using the current dual iterate $(1, y)$ as the dual certificate, a generic primal pair $(\gamma, -b)$, where γ is an unknown scalar parameter.

We can then proceed in one of two ways: for some cones, such as the nonnegative orthant and the positive semidefinite cone, it is possible to find the best certifiable γ directly from this affine expression. (For linear programming, this is simple arithmetic, for semidefinite programming it leads to a generalized eigenvalue problem.)

If this is not possible or efficient, we may invoke the sufficient conditions found in [Section 4](#). Comparing [Eq. \(6.5\)](#) to those results, we see that every vector (y_0, y) in a suitable Dikin ellipsoid centered at $(1, y_\tau)$ is a dual certificate of the relation $(b^\top y_\tau + \tau\nu, -b) \in \mathcal{K}_{c,A}^*$, certifying that there exists a feasible solution x in [\(5.2\)](#) whose attained value is $c^\top x = b^\top y_\tau + \tau\nu$. Conversely, $(1, y)$ is a dual certificate of $-g(1, y)$ and also of every vector in the Dikin ellipsoid of radius r around this center. (Here, $r = 1$ for H-certificates and $r = (2(\nu + 1))^{-1}$ for B-certificates, based on [Theorems 4.1](#) and [4.2](#).) Thus, for every y sufficiently close to y_τ , we have that

$$(6.6) \quad (\gamma, -b) \in \mathcal{B}^*(-g(1, y), r)$$

for some γ . With y and the radius r fixed, the set of values γ for which [\(6.6\)](#) holds is an interval that can be easily computed by solving a scalar quadratic equation. (To set up this equation, we need a factorization or the inverse of the Hessian of $f_{\mathcal{K}_{c,A}^*}$ at y , but this is already necessary to run the IPAs. The cost of solving this equation is therefore negligible.) What our arguments above show is that *in a (known) neighborhood of the central path, the smallest γ value in this interval is no greater than $b^\top y_\tau + \tau\nu$, and therefore the current dual solution and the reconstructed primal solution will have a small duality gap.*

EXAMPLE 6.2. Consider a standard form linear program and its dual, [Eq. \(5.2\)](#) with

$$\mathcal{K} = \mathbb{R}_+^4, \quad c^\top = (1, 1, 1, 1), \quad A = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 0 & -6 & 0 & 1 \end{pmatrix}, \quad b = \begin{pmatrix} 19 \\ -5 \end{pmatrix}.$$

The optimal objective function value is 5.5, attained at $x = (0, 1.5, 0, 4)^\top$. A crudely calculated near-optimal dual solution near $y_{0.1}$, the point corresponding to $\tau = 0.1$ on the central path, is

$$y = (6/23, -2/29)^\top \quad \text{with} \quad b^\top y = 3536/667 \approx 5.30135.$$

²It is possible to make this statement and the remaining “for sufficiently close...” statements of this section more precise and quantitative using [Theorem 4.1](#) and [Theorem 4.2](#), but we omit the details, as some path-following methods may use different neighborhoods anyway.

We can plug this dual solution (with $y_0 = 1$) and our barrier (6.2) into the reconstruction formula with the target primal objective function value γ . The resulting primal solution candidate is

$$x_\gamma = \begin{pmatrix} \frac{10459868\gamma - 57626369}{13715262} \\ \frac{21391355 - 2022605\gamma}{6857631} \\ \frac{31571864\gamma - 173276291}{13715262} \\ \frac{-4045210\gamma + 31353325}{2285877} \end{pmatrix},$$

and it can be directly verified that $c^T x_\gamma = \gamma$ and $Ax_\gamma = b$ independent of γ , as desired. In this simple example, rather than using the sufficient condition (6.6), we can directly compute the smallest γ for which $x_\gamma \geq 0$. The best attainable primal objective value thus certified is

$$\gamma = 57626369/10459868 \approx 5.50928.$$

Note that, perhaps surprisingly, the certified primal objective is considerably closer to optimal than the value attained by the dual feasible solution used to certify it.

7. Application to polynomial optimization. In this section, we explore the application of H- and B-certificates in the context of certificates proving the nonnegativity of polynomials. As it will be apparent, our tools developed in this paper can be directly applied to all the tractable nonnegativity certificates popular in the polynomial optimization literature.

Recall that a polynomial optimization problem (POP) is a problem of the form

$$(7.1) \quad \inf \{p(x) \mid q_i(x) \geq 0 \ (i = 1, \dots, m)\},$$

where p and q_1, \dots, q_m are n -variate polynomials. This, generally intractable, non-convex optimization problem is often approached as a (convex) optimization problem over the cone of nonnegative polynomials:

$$\sup \{\alpha \in \mathbb{R} \mid p - \alpha \in \mathcal{P}_\Delta\},$$

where \mathcal{P}_Δ is the cone of polynomials that are nonnegative over the set

$$(7.2) \quad \Delta \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid q_i(x) \geq 0 \ (i = 1, \dots, m)\}.$$

The strategy then is to replace \mathcal{P}_Δ with a convex subcone $\mathcal{C} \subseteq \mathcal{P}_\Delta$ whose membership problem is tractable, such as a (weighted) sums-of-squares (SOS) or sums of nonnegative circuit (SONC) polynomial cone:

$$(7.3) \quad \sup \{\alpha \in \mathbb{R} \mid p - \alpha \in \mathcal{C}\}.$$

Every feasible solution of (7.3) yields a lower bound on the optimal value of (7.1). This can be computed as a heuristic (with no expectation that the resulting lower bound is close to the infimum of the original problem), or as part of a scheme (theoretically supported by a *Positivstellensatz*, i.e., a characterization of positive polynomials, such as [14, 33, 10, 19]), where a sequence of increasingly good approximations in the form of nested cones

$$\mathcal{C}_1 \subseteq \mathcal{C}_2 \subseteq \dots \quad \text{satisfying} \quad \text{cl}(\cup_{i=1}^\infty \mathcal{C}_i) = \mathcal{P}_\Delta$$

is available, leading to a sequence of lower bounds converging to the infimum of (7.1).

In the presence of strong duality, if we are only concerned about the approximate numerical value of this lower bound, we may opt to solve the dual problem of (7.3) instead; for example, if \mathcal{C} is a suitable SOS cone, (7.3) is known as the *SOS relaxation* of the POP (7.1), and the dual of (7.3) is known as *Lasserre’s moment relaxation* [20].

The theoretical conundrum is now similar to the issue we have encountered in the Introduction and Section 6: every cone \mathcal{C} that is typically used in this context to approximate \mathcal{P}_Δ , including each of those examined later in this section, is represented as a linear image \mathcal{K}_A of a “nice” cone \mathcal{K} (recall Eq. (3.1)), which means that a typical numerical method solving (7.3) will result only in an *approximate* certificate, rigorously proving that some polynomial *close to, but not identical to* $p - \alpha$ is nonnegative over Δ [21]. On the dual side, a near-optimal feasible solution certifiably in $(\mathcal{C}^*)^\circ$ can be found, but this only yields an *upper bound* on the optimal value of (7.3), that is, an upper bound on a lower bound on the optimal value of (7.1). In summary, both (7.3) and its dual may be solvable approximately using numerical methods, but neither numerical solutions provide a theoretically satisfactory, rigorously certified lower bound (with a nonnegativity certificate that is verifiable in rational arithmetic) on the optimal value of (7.1). Instead, additional post-processing, in the form of some rounding procedure, is required to turn the numerical bounds and certificates into rigorous ones, which (in addition to being additional work) puts additional constraints on the quality of the numerical solution. See, e.g., [31, 16, 21] for examples in the context of SOS certification.

The dual certificates described in this paper, and the reconstruction formula (3.5) resolve this problem in a simple, unifying framework, by allowing us to compute explicit, exact SOS, SONC, etc. representations, that is, near-optimal (exactly) feasible solutions to (7.3), from suitable near-optimal solutions of the dual problem for various cones \mathcal{C} . As an additional theoretical benefit, dual certificates provide “parsimonious” characterizations of polynomials in \mathcal{C} , in contrast with the usual representations of these cones as images of much higher dimensional cones. (Although, if necessary, those representations can also be constructed in closed form using the reconstruction formula.)

In the following subsections, we discuss these characterizations for the most commonly used tractable subcones of nonnegative polynomials. For ease of presentation, we focus on nonnegativity over $\Delta = \mathbb{R}^n$; the extension to the general case is straightforward as long as an appropriate Positivstellensatz is available. For instance, when Δ is compact, the results of the recent work [19] cover at least one appropriate Positivstellensatz using every technique considered in this section: SOS, SONC, DSOS, SDSOS and SAGE polynomials, and nonnegative AG functions. Throughout, we shall represent polynomials in the monomial basis, using the common shorthand $x^\alpha \stackrel{\text{def}}{=} \prod_{i=1}^n x_i^{\alpha_i}$ to represent monomials.

7.1. Sums-of-squares certificates. We return to Example 2.4 to develop some additional detail regarding H-certificates in sums-of-squares cones.

DEFINITION 7.1. *An n -variate polynomial of degree $2d$ is sums-of-squares (SOS) if it can be written as a sum of squares of degree- d polynomials. The set of such polynomials is denoted by $\Sigma_{n,d}$.*

The set $\Sigma_{n,d} \subseteq \mathbb{R}^N$ with $N = \binom{n+2d}{n}$ is a proper convex cone with a well-known semidefinite representation, briefly recalled for completeness in Proposition 7.2 below. We need the following notation: \mathbb{S}^M denotes the set of real, symmetric matrices of order M , and \mathbb{S}_+^M denotes the cone of positive semidefinite matrices within \mathbb{S}^M . If $A(\cdot)$ is a linear map between finite dimensional real linear spaces, $A^*(\cdot)$ denotes its adjoint.

PROPOSITION 7.2. *Given positive integers d and n , denote by $m_d(x)$ the vector of monomials of degree up to d in the variables $x = (x_1, \dots, x_n)$, let $M = \binom{n+d}{n}$ and $N = \binom{n+2d}{n}$. A polynomial $p : x \mapsto \sum_{i=1}^N p_{\alpha_i} x^{\alpha_i}$ is SOS if and only if*

$$(7.4) \quad p(x) = m_d(x)^T X m_d(x)$$

identically for some $X \in \mathbb{S}_+^M$. In other words, there exists a linear map $A : \mathbb{S}^M \rightarrow \mathbb{R}^N$ such that

$$p \in \Sigma_{n,d} \iff \exists X \in \mathbb{S}_+^M : p = A(X).$$

The dual-SOS cone, often called the pseudo-moment cone, in turn, also has a semidefinite representation:

$$y \in \Sigma_{n,d}^* \iff A^*(y) \in \mathbb{S}_+^M.$$

Thus, $\Sigma_{n,d}$ is clearly of the form \mathcal{K}_A in (3.1), with \mathbb{S}_+^M playing the role of \mathcal{K} . Note that in general, $N \gg M$, that is, this representation characterizes the SOS cone as the linear image of a much higher-dimensional semidefinite cone. The linear map A in Proposition 7.2 can be easily constructed; in matrix form it is a zero-one matrix that encodes which monomials can be obtained as a product of pairs of monomials.

Analogous statements can be formulated for polynomials represented in other bases, such as the Chebyshev basis of the first or second kind, which helps numerical conditioning, and also for weighted SOS polynomials nonnegative over semialgebraic sets Δ ; this requires no changes other than replacing A^* with another suitable, easy-to-construct, linear map Λ that depends on the bases and the polynomials representing Δ in Eq. (7.2); see, e.g., [22].

It is well-known that \mathbb{S}_+^M is the hyperbolicity cone associated with the determinant function (a hyperbolic polynomial), and as a result, the function $f : X \mapsto -\log \det(X)$ is an M -LHSCB for this cone [12]. Consequently, $\Sigma_{n,d}^*$ is also a hyperbolicity cone with the corresponding M -LHSCB $f_{\Sigma_{n,d}^*} : y \mapsto -\log \det \Lambda(y)$, and H-certificates defined using the Hessian of this function are rigorous nonnegativity certificates. Often, the operator Λ corresponding to the bases and weights used to represent the cone maps rational vectors to rational matrices. This means that numerically computed H-certificates are automatically rational nonnegativity certificates that can be verified in polynomial time (in the bit model) in rational arithmetic [9].

7.2. SONC certificates.

DEFINITION 7.3. *We say that an n -variate polynomial p is a circuit polynomial if its support can be written as $\text{supp}(p) = \{\alpha_1, \dots, \alpha_r, \beta\} \subset \mathbb{N}^n$, where the set $\{\alpha_1, \dots, \alpha_r\}$ is affinely independent and $\beta = \sum_{i=1}^r \lambda_i \alpha_i$ with some $\lambda_i > 0$ satisfying $\sum_{i=1}^r \lambda_i = 1$. In other words, β lies in the convex hull of the α_i , and the scalars λ_i are the corresponding (uniquely defined, strictly positive) barycentric coordinates of β . The support set of a circuit polynomial is called a circuit.*

The cone of nonnegative circuit polynomials is closely related to the notion of power cones. The literature has conflicting notation and terminology; we shall use the following definitions, compatible with our earlier Example 2.2:

DEFINITION 7.4. *The (generalized) power cone with signature $\lambda = (\lambda_1, \dots, \lambda_r)$ is the convex cone defined as*

$$\mathcal{P}_\lambda \stackrel{\text{def}}{=} \{(v, z) \in \mathbb{R}_+^r \times \mathbb{R} \mid |z| \leq v^\lambda\}.$$

Its dual cone is the (generalized) dual power cone:

$$\mathcal{P}_\lambda^* \stackrel{\text{def}}{=} \left\{ (v, z) \in \mathbb{R}_+^r \times \mathbb{R} \mid |z| \leq \prod_{i=1}^r \left(\frac{v_i}{\lambda_i} \right)^{\lambda_i} \right\}.$$

It can be shown that \mathcal{P}_λ and \mathcal{P}_λ^* are proper convex cones for every signature $\lambda \in (0, 1)^r$ [4, 11]. With this notation, it is easy to characterize the subset of circuit polynomials nonnegative on \mathbb{R}^n :

PROPOSITION 7.5 ([15, 26]). *Let p be an n -variate circuit polynomial satisfying $p(x) = \sum_{i=1}^r p_{\alpha_i} x^{\alpha_i} + p_\beta x^\beta$ for some real coefficients p_{α_i} and p_β , and suppose that $\beta = \sum_{i=1}^r \lambda_i \alpha_i$ with some $\lambda_i > 0$ satisfying $\sum_{i=1}^r \lambda_i = 1$. Then p is nonnegative if and only if $\alpha_i \in (2\mathbb{N})^n$ and $p_{\alpha_i} > 0$ for each i , and at least one of the following two alternatives holds:*

1. $\beta \in (2\mathbb{N})^n$ and $p_\beta \geq 0$, or
2. $((p_{\alpha_1}, \dots, p_{\alpha_r}), p_\beta) \in \mathcal{P}_\lambda^*$.

Proposition 7.5 has a simple interpretation: a circuit polynomial is nonnegative if and only if it is either a sum of nonnegative monomials, also known as *monomial squares*, or if its nonnegativity can be directly proven by applying the weighted AM/GM inequality to its monomials, with weights λ .

DEFINITION 7.6. *We say that a polynomial is a sum of nonnegative circuit polynomials, or SONC for short, if it can be written as a sum of monomial squares and nonnegative circuit polynomials. An explicit representation of a polynomial in this form is called a SONC decomposition.*

We are ready to state the key representation theorem of SONC polynomials.

PROPOSITION 7.7 ([26, Theorem 5]). *Every SONC polynomial p has a SONC decomposition in which every nonnegative circuit polynomial and monomial square is supported on a subset of $\text{supp}(p)$.*

Thus, in the notation of Section 3, the cone of SONC polynomials of degree $\deg(p)$ supported on $\text{supp}(p)$ is a cone of the form $\mathcal{K}_A \subseteq \mathbb{R}^{|\text{supp}(p)|}$, where the high-dimensional cone \mathcal{K} is the Cartesian product of dual power cones (one for each circuit supported on $\text{supp}(p)$) and nonnegative half-lines (one for each monomial square). The matrix A of the linear transformation is a matrix of zeros and ones that tallies the coefficients of each monomial from the different circuits and monomial squares they appear in.

Since the dual cone \mathcal{K}^* is the Cartesian product of a nonnegative orthant and power cones of the form \mathcal{P}_λ with different signatures λ , in order to apply the techniques of Section 3, we only need an LHSCB for \mathcal{P}_λ . Such barriers have already been found by a number of authors [4, 35]. For example,

$$f_{\mathcal{P}_\lambda}(v, z) = -\ln(v^{2\lambda} - z^2) - \sum_{i=1}^n (1 - \lambda_i) \ln(v_i)$$

is an $(n+1)$ -LHSCB for this cone [35]. Note, however, that \mathcal{P}_λ is not a hyperbolicity cone for every λ , and we have already shown in [Example 2.2](#) that H-certificates do not work for, e.g., $\lambda = (2/3, 1/3)$. Therefore, at least for general λ , B-certificates are needed for rigorous nonnegativity certification.

7.3. Signomial functions and S-cones. The concept of SONC polynomials can be straightforwardly generalized to nonnegative n -variate functions of the form

$$(7.5) \quad p(x) = \sum_{\alpha \in \mathcal{A}} c_\alpha |x|^\alpha + \sum_{\beta \in \mathcal{B}} d_\beta x^\beta$$

with exponents $\mathcal{A} \subseteq \mathbb{R}^n$ and $\mathcal{B} \subseteq \mathbb{N}^n \setminus (2\mathbb{N})^n$, using $|x|$ to denote component-wise absolute value. The exponents in \mathcal{A} need not be integers, so this set encapsulates not only polynomials of x and $|x|$, but also signomial functions [3]. Analogously to circuit polynomials in the SONC setting, we define “atomic” functions in this set, called AG functions, for which nonnegativity holds if and only if it can be proven via a single application of the weighted AM/GM inequality:

DEFINITION 7.8. *Let p be a function of the form (7.5). We say that*

1. *p is an even AG function with support $(\mathcal{A}, \mathcal{B})$ if at most one c_α is negative and all d_β are zero; and*
2. *p is an odd AG function with support $(\mathcal{A}, \mathcal{B})$ if all c_α are nonnegative and at most one d_β is nonzero.*

In addition, p is called an AG function with support $(\mathcal{A}, \mathcal{B})$ if p is an even AG function or an odd AG function with the same support. Finally, the S-cone $\mathcal{S}_{\mathcal{A}, \mathcal{B}}$ is the set of AG functions nonnegative on \mathbb{R}^n .

Katthän, Naumann and Theobald have shown that, analogously to the SONC results, functions in $\mathcal{S}_{\mathcal{A}, \mathcal{B}}$ can be written as sums of AG functions with support $(\mathcal{A}, \mathcal{B})$ [17]. They have also characterized both even and odd AG functions in terms of at most two relative entropy cone inequalities [17, Theorem 2.8]. This representation is compatible with the concept of dual certificates. For example, an n -variate odd AG function given by (7.5) with support $(\mathcal{A}, \{\beta\})$ is nonnegative on \mathbb{R}^n if and only if there exists a $\nu \in \mathbb{R}_+^{\mathcal{A}}$ such that

$$\sum_{\alpha \in \mathcal{A}} \nu_\alpha \alpha = \left(\sum_{\alpha \in \mathcal{A}} \nu_\alpha \right) \beta, \quad (d, \nu, ec) \in \mathcal{R}_{|\mathcal{A}|}, \quad \text{and} \quad (-d, \nu, ec) \in \mathcal{R}_{|\mathcal{A}|},$$

and where $e = \exp(1)$ and \mathcal{R}_N represents the (vector) relative entropy cone

$$(7.6) \quad \mathcal{R}_N \stackrel{\text{def}}{=} \text{cl} \left\{ (u, v, w) \in \mathbb{R} \times \mathbb{R}_{++}^N \times \mathbb{R}_{++}^N \mid u \geq \sum_{i=1}^N v_i \ln \left(\frac{v_i}{w_i} \right) \right\}.$$

A similar characterization can be given for nonnegative even AG functions. As a result, membership in $\mathcal{S}_{\mathcal{A}, \mathcal{B}}$ can be B-certified using any known LHSCB for the dual of the relative entropy cone, \mathcal{R}_N^* . This cone is also already well-studied, and it also has a suitable barrier function:

PROPOSITION 7.9 ([5]). *The dual of the cone \mathcal{R}_N given in Eq. (7.6) is the cone*

$$\mathcal{R}_N^* \stackrel{\text{def}}{=} \text{cl} \left\{ (u, v, w) \in \mathbb{R}_{++} \times \mathbb{R}_{++}^N \times \mathbb{R}^N \mid w_i \geq u \left(\ln \left(\frac{u}{v_i} \right) - 1 \right) \quad (i = 1, \dots, N) \right\}.$$

The function

$$f_{\mathcal{R}_N^*}(u, v, w) \stackrel{\text{def}}{=} - \sum_{i=1}^N \ln \left(w_i - u \ln \left(\frac{u}{v_i} \right) + u \right) - N \ln(u) - \sum_{i=1}^N \ln(v_i)$$

is a $(3N)$ -LHSCB for \mathcal{R}_N^ .*

Once again, B-certificates are required, because (similarly to the closely related exponential cone and its dual), relative entropy cones and their duals do not satisfy the relation (2.2).

7.4. Nonnegative bases, LP- and SOCP-representable subcones of SOS polynomials. In an attempt to get around the high complexity of solving large-scaled SDPs associated with SOS certificates (Proposition 7.2), various authors have suggested using polyhedral or second-order cone representable subcones of SOS polynomials [14, 18, 19]. These include the following, in increasing order of complexity:

1. Polynomials with nonnegative coefficients in a fixed basis chosen to be trivially nonnegative on Δ ;
2. Polynomials of the form (7.4) with a diagonally dominant X whose diagonal is componentwise nonnegative (sometimes referred to as DSOS polynomials);
3. Polynomials of the form (7.4), where X is a matrix with factor-width no greater than k ; i.e., X is a sum of PSD matrices that are non-zero only in a single principal submatrix of order k [2]. (In the $k = 2$ case, these are sometimes called SDSOS polynomials.)

It is clear that each of these cones are of the form \mathcal{K}_A where \mathcal{K} is either the nonnegative orthant or a Cartesian product of PSD matrices of order at most k . In Cases (1) and (3) this is immediate from the definition. (Also recall that the cone of 2×2 PSD matrices is linearly isomorphic to a three-dimensional second-order cone, making the SDSOS cone SOCP-representable.) In Case (2) we refer to the fact that diagonally dominant positive semidefinite matrices of order n have n^2 explicitly known extreme rays, and every matrix in the cone can be written as a conic combination of these extreme vectors; see, e.g., [1, Theorem 2] or [7, Proposition 1]. Therefore, H-certificates obtained from the standard logarithmic barrier of \mathcal{K}_A^* can be used to rigorously certify membership in this cone.

As with SOS polynomials, $\dim(\mathcal{K}) \gg \dim(\mathcal{K}_A)$ (particularly in Cases (1) and (3) above), and the matrix A is an easily computable matrix with rational entries in the case of DSOS and SDSOS polynomials. It is usually a rational matrix in Case (1), too, unless the chosen basis polynomials have irrational coefficients. Thus, numerically computed H-certificates (see Section 5) are automatically rational vectors whose correctness can be verified in rational arithmetic if the certified polynomial $p - \alpha$ has rational coefficients.

8. Discussion. We presented a theoretical framework of cone membership certificates that allows us to interpret vectors in the dual cone as certificates of membership in the primal cone. The new certification scheme builds on the classic self-concordance theory of Nesterov and Nemirovski [25], but it goes beyond the well-known result that the gradient map associated with an SCB defines a bijection between the interior of a cone and the interior of its dual. The key difference is that $\mathcal{H}(b)$ and $\mathcal{B}(b)$ are full-dimensional cones, which means that we can use numerical methods to identify any of the certificates, and still establish rigorous certification of membership. Similarly, in the case of a low-dimensional projection \mathcal{K}_A of the high-dimensional cone \mathcal{K} , an exact preimage $x \in \mathcal{K}$ of a vector $b \in \mathcal{K}_A$ can be computed from any (numerically computed) dual certificate. The same infrastructure also allowed us to compute exact primal feasible solutions of general convex conic optimization problems from suitable dual feasible solutions, with duality gaps that approach zero as the dual feasible solution approaches optimality.

We have also established that for many cones, some already existing optimization methods (specifically, feasible dual and primal-dual IPAs operating in appropriate neighborhoods of the central path), can be used to compute dual certificates. For certain cones (e.g., linear images of symmetric cones), the existing dual methods automatically compute dual certificates, and

these can now be extended to compute exact primal feasible solutions that converge to a primal optimal solution, complementing the dual feasible iterates of the method. For the most general case of nonsymmetric conic optimization, the IPA recently proposed by the authors [27] can be applied.

In the future, we expect that this theory will lead to more efficient optimization methods for certain structured optimization problems. As outlined in the paper, optimization over a cone \mathcal{K}_A can now circumvent the conventional “lifting” to optimization over \mathcal{K} (also known as the “extended formulation”) for a potential reduction in the solution time and especially in the required working memory. As long as \mathcal{K}^* admits an LHSCB, we can directly optimize over \mathcal{K}_A^* with the inherited barrier and reconstruct the desired primal solution in \mathcal{K}_A after the fact.

Looking beyond optimization, the exactness of the reconstruction formula may be used to establish rigorous computer-generated proofs based on numerical optimization methods via rigorous nonnegativity certificates of polynomials. (E.g., exact SOS certificates of global lower bounds of a POP can be constructed from suitable feasible solutions of the Lasserre moment relaxation.) The last section of our paper lays the necessary theoretical groundwork for this; a more detailed study is the subject of a forthcoming paper.

REFERENCES

- [1] G. BARKER AND D. CARLSON, *Cones of diagonally dominant matrices*, Pacific Journal of Mathematics, 57 (1975), pp. 15–32, <https://doi.org/10.2140/pjm.1975.57.15>.
- [2] G. BLEKHERMAN, S. S. DEY, K. SHU, AND S. SUN, *Hyperbolic relaxation of k -locally positive semidefinite matrices*, SIAM Journal on Optimization, 32 (2022), pp. 470–490, <https://doi.org/10.1137/20M1387407>.
- [3] V. CHANDRASEKARAN AND P. SHAH, *Relative entropy relaxations for signomial optimization*, SIAM Journal on Optimization, 26 (2016), pp. 1147–1173, <https://doi.org/10.1137/140988978>.
- [4] R. CHARES, *Cones and interior-point algorithms for structured convex optimization involving powers and exponentials*, PhD thesis, Université Catholique de Louvain, 2009.
- [5] Y. CHEN AND P. GOULART, *An efficient IPM implementation for a class of nonsymmetric cones*, arXiv:2305.12275, (2023), <https://arxiv.org/abs/2305.12275>.
- [6] C. COEY, L. KAPELEVICH, AND J. P. VIELMA, *Performance enhancements for a generic conic interior point algorithm*, Mathematical Programming Computation, 15 (2023), pp. 53–101, <https://doi.org/10.1007/s12532-022-00226-0>.
- [7] G. DAHL, *A note on diagonally dominant matrices*, Linear Algebra and its Applications, 317 (2000), pp. 217–224, [https://doi.org/10.1016/S0024-3795\(00\)00178-6](https://doi.org/10.1016/S0024-3795(00)00178-6).
- [8] M. M. DAVIS AND D. PAPP, *Dual certificates and efficient rational sum-of-squares decompositions for polynomial optimization over compact sets*, SIAM Journal on Optimization, 32 (2022), pp. 2461–2492, <https://doi.org/10.1137/21M1422574>.
- [9] M. M. DAVIS AND D. PAPP, *Rational dual certificates for weighted sums-of-squares polynomials with boundable bit size*, Journal of Symbolic Computation, 121 (2024), p. 102254, <https://doi.org/10.1016/j.jsc.2023.102254>.
- [10] M. DRESSLER, S. ILIMAN, AND T. DE WOLFF, *A Positivstellensatz for sums of nonnegative circuit polynomials*, SIAM Journal on Applied Algebra and Geometry, 1 (2017), pp. 536–555, <https://doi.org/10.1137/16M1086303>.
- [11] M. DRESSLER, H. NAUMANN, AND T. THEOBALD, *The dual cone of sums of non-negative circuit polynomials*, Advances in Geometry, 21 (2021), pp. 227–236, <https://doi.org/10.1515/advgeom-2020-0019>.
- [12] O. GÜLER, *Hyperbolic polynomials and interior point methods for convex programming*, Mathematics of Operations Research, 22 (1997), pp. 350–377, <https://doi.org/10.1287/moor.22.2.350>.
- [13] O. GÜLER, *Foundations of optimization*, vol. 258 of Graduate Texts in Mathematics, Springer, New York, NY, 2010, <https://doi.org/10.1007/978-0-387-68407-9>.
- [14] D. HANDELMAN, *Representing polynomials by positive linear functions on compact convex polyhedra*, Pacific Journal of Mathematics, 132 (1988), pp. 35–62, <https://doi.org/10.2140/pjm.1988.132.35>.
- [15] S. ILIMAN AND T. DE WOLFF, *Amoebas, nonnegative polynomials and sums of squares supported on circuits*, Research in the Mathematical Sciences, 3 (2016), pp. 1–35, <https://doi.org/10.1186/>

- s40687-016-0052-2.
- [16] E. L. KALTOFEN, B. LI, Z. YANG, AND L. ZHI, *Exact certification in global polynomial optimization via sums-of-squares of rational functions with rational coefficients*, Journal of Symbolic Computation, 47 (2012), pp. 1–15, <https://doi.org/http://dx.doi.org/10.1016/j.jsc.2011.08.002>.
 - [17] L. KATTHÄN, H. NAUMANN, AND T. THEOBALD, *A unified framework of SAGE and SONC polynomials and its duality theory*, Mathematics of Computation, 90 (2021), pp. 1297–1322, <https://doi.org/10.1090/mcom/3607>.
 - [18] X. KUANG, B. GHADDAR, J. NAOUM-SAWAYA, AND L. F. ZULUAGA, *Alternative SDP and SOCP approximations for polynomial optimization*, EURO Journal on Computational Optimization, 7 (2019), pp. 153–175, <https://doi.org/10.1007/s13675-018-0101-2>.
 - [19] O. KURYATNIKOVA, J. C. VERA, AND L. F. ZULUAGA, *Reducing nonnegativity over general semialgebraic sets to nonnegativity over simple sets*, SIAM Journal on Optimization, 34 (2024), pp. 1970–2006, <https://doi.org/10.1137/22M1501027>.
 - [20] J. B. LASSERRE, *Global optimization with polynomials and the problem of moments*, SIAM Journal on Optimization, 11 (2001), pp. 796–817, <https://doi.org/10.1137/S1052623400366802>.
 - [21] V. MAGRON AND M. S. E. DIN, *RealCertify: a Maple package for certifying non-negativity*, ACM Communications in Computer Algebra, 52 (2018), pp. 34–37, <https://doi.org/10.1145/3282678.3282681>.
 - [22] Y. NESTEROV, *Squared functional systems and optimization problems*, in High Performance Optimization, H. Frenk, K. Roos, T. Terlaky, and S. Zhang, eds., Kluwer, Dordrecht, The Netherlands, 2000, pp. 405–440, https://doi.org/10.1007/978-1-4757-3216-0_17.
 - [23] Y. NESTEROV, *Constructing self-concordant barriers for convex cones*, Tech. Report CORE 2006/30, Center for Operations Research and Econometrics (CORE), Catholic University of Louvain (UCL), Belgium, Mar. 2006.
 - [24] Y. NESTEROV, *Lectures on convex optimization*, vol. 137 of Springer Optimization and Its Applications, Springer Nature, Cham, Switzerland, 2nd ed., 2018, <https://doi.org/10.1007/978-3-319-91578-4>.
 - [25] Y. NESTEROV AND A. NEMIROVSKII, *Interior-point polynomial algorithms in convex programming*, vol. 13 of SIAM Studies in Applied Mathematics, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1994, <https://doi.org/10.1137/1.9781611970791>.
 - [26] D. PAPP, *Duality of sum of nonnegative circuit polynomials and optimal SONC bounds*, Journal of Symbolic Computation, 114 (2023), pp. 246–266, <https://doi.org/10.1016/j.jsc.2022.04.015>.
 - [27] D. PAPP AND A. VARGA, *Interior-point algorithms with full Newton steps for nonsymmetric convex conic optimization*, arXiv:2502.16020, (2025), <https://arxiv.org/abs/2502.16020>.
 - [28] D. PAPP AND S. YILDIZ, *Sum-of-squares optimization without semidefinite programming*, SIAM Journal on Optimization, 29 (2019), pp. 822–851, <https://doi.org/10.1137/17M1160124>.
 - [29] D. PAPP AND S. YILDIZ, *Alfonso: Matlab package for nonsymmetric conic optimization*, INFORMS Journal on Computing, 34 (2022), pp. 11–19, <https://doi.org/10.1287/ijoc.2021.1058>.
 - [30] P. A. PARRILO, *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*, PhD thesis, California Institute of Technology, May 2000, <https://doi.org/10.7907/2K6Y-CH43>.
 - [31] H. PEYRL AND P. A. PARRILO, *Computing sum of squares decompositions with rational coefficients*, Theoretical Computer Science, 409 (2008), pp. 269–281, <https://doi.org/10.1016/j.tcs.2008.09.025>.
 - [32] V. POWERS, *Rational certificates of positivity on compact semialgebraic sets*, Pacific Journal of Mathematics, 251 (2011), pp. 385–391, <https://doi.org/10.2140/pjm.2011.251.385>.
 - [33] M. PUTINAR, *Positive polynomials on compact semi-algebraic sets*, Indiana University Mathematics Journal, 42 (1993), pp. 969–984.
 - [34] J. RENEGAR, *A mathematical view of interior-point methods in convex optimization*, MOS-SIAM Series on Optimization, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001, <https://doi.org/10.1137/1.9780898718812>.
 - [35] S. ROY AND L. XIAO, *On self-concordant barriers for generalized power cones*, Optimization Letters, 16 (2022), pp. 681–694, <https://doi.org/10.1007/s11590-021-01748-7>.
 - [36] S. A. SERRANO, *Algorithms for unsymmetric cone optimization and an implementation for problems with the exponential cone*, Stanford University, 2015.
 - [37] A. SKAJAA AND Y. YE, *A homogeneous interior-point algorithm for nonsymmetric convex conic optimization*, Mathematical Programming, 150 (2015), pp. 391–422, <https://doi.org/10.1007/s10107-014-0773-1>.