

SUBLINEAR CIRCUITS AND THE CONSTRAINED SIGNOMIAL NONNEGATIVITY PROBLEM

RILEY MURRAY, HELEN NAUMANN, AND THORSTEN THEOBALD

ABSTRACT. Conditional Sums-of-AM/GM-Exponentials (conditional SAGE) is a decomposition method to prove nonnegativity of a signomial or polynomial over some subset X of real space. In this article, we undertake the first structural analysis of conditional SAGE signomials for convex sets X . We introduce the X -circuits of a finite subset $\mathcal{A} \subset \mathbb{R}^n$, which generalize the simplicial circuits of the affine-linear matroid induced by \mathcal{A} to a constrained setting. The X -circuits serve as the main tool in our analysis and exhibit particularly rich combinatorial properties for polyhedral X , in which case the set of X -circuits is comprised of one-dimensional cones of suitable polyhedral fans.

The framework of X -circuits transparently reveals when an X -nonnegative conditional AM/GM-exponential can in fact be further decomposed as a sum of simpler X -nonnegative signomials. We develop a duality theory for X -circuits with connections to geometry of sets that are convex according to the geometric mean. This theory provides an optimal power cone reconstruction of conditional SAGE signomials when X is polyhedral. In conjunction with a notion of reduced X -circuits, the duality theory facilitates a characterization of the extreme rays of conditional SAGE cones.

Since signomials under logarithmic variable substitutions give polynomials, our results also have implications for nonnegative polynomials and polynomial optimization.

1. INTRODUCTION

Given a finite subset $\mathcal{A} \subset \mathbb{R}^n$, a signomial on \mathcal{A} is a real-linear combination

$$f = \sum_{\alpha \in \mathcal{A}} c_{\alpha} e^{\alpha} \text{ of basis functions } e^{\alpha}(x) := \exp(\alpha^T x) \quad (1)$$

with coefficients $c = (c_{\alpha})_{\alpha \in \mathcal{A}}$. Signomials are a fundamental class of functions with applications in, but not limited to, chemical reaction networks [19, 20], aircraft design optimization [31, 42], and epidemiological process control [29, 35]. We refer the reader to [10] and its references for the manifold occurrences of signomials in pure and applied mathematics and to [22, §1.1] for an abridged history of signomial modeling which begins with geometric programming. Signomials are often considered under a logarithmic change of variables $y \mapsto f(\log y) = \sum_{\alpha \in \mathcal{A}} c_{\alpha} \prod_{i=1}^n y_i^{\alpha_i}$, so that for $\mathcal{A} \subset \mathbb{N}^n$ one obtains polynomials over the positive orthant \mathbb{R}_{++}^n .

A basic question one might ask of a signomial is when the coefficients $c = (c_{\alpha})_{\alpha \in \mathcal{A}}$ are such that f is globally nonnegative. Framing this question in terms of a signomial's coefficients

Date: January 20, 2022.

2010 Mathematics Subject Classification. 14P05, 90C23, 90C30 (primary), 05B35, 52A20 (secondary).

Key words and phrases. Sums of arithmetic-geometric exponentials, positive signomials, exponential sums, sums of nonnegative circuit polynomials (SONC), positive polynomials, multiplicative convexity, log convex sets.

affords direct connections to polynomials. If the exponent vectors \mathcal{A} are contained in \mathbb{N}^n , then f is nonnegative on \mathbb{R}^n if and only if the polynomial $y \mapsto \sum_{\alpha \in \mathcal{A}} c_\alpha \prod_{i=1}^n y_i^{\alpha_i}$ is nonnegative on the nonnegative orthant \mathbb{R}_+^n . Deciding such nonnegativity problems is NP-hard in general [25]. However, several researchers have developed sufficient conditions for nonnegativity based on the AM/GM inequality. In contrast to the well-known Sums-of-Squares nonnegativity certificates in the polynomial setting (see, e.g., [18, 34]), the techniques based on the AM/GM inequality are not tied to the notion of a polynomial's degree, and hence apply to general signomials. The earliest results here are due to Reznick [36], with a resurgence marked by the works of Pantea, Koepl, and Craciun [32], Ilman and de Wolff [14], and Chandrasekaran and Shah [4]. Whether considered for signomials or polynomials, such techniques have appealing forms of sparsity preservation in the proofs of nonnegativity [23, 40].

In this article, we are concerned with the question of when a signomial on exponents \mathcal{A} is X -nonnegative (i.e., nonnegative on X) for a convex set X . In important progress on this question, Murray, Chandrasekaran and Wierman have proposed an extension of the *Sums-of-AM/GM-Exponentials* or *SAGE* approach to global nonnegativity [4], which goes by the name *conditional SAGE* [24]. The method works as follows: if a signomial f of the form (1) has at most one negative coefficient c_β , i.e.,

$$f = \sum_{\alpha \in \mathcal{A} \setminus \{\beta\}} c_\alpha e^\alpha + c_\beta e^\beta \quad \text{with} \quad c_\alpha \geq 0 \quad \text{for all } \alpha \in \mathcal{A} \setminus \{\beta\}$$

then we may divide out the corresponding basis function e^β to obtain a new signomial $g = \sum_{\alpha \in \mathcal{A}} c_\alpha e^{[\alpha - \beta]}$ without affecting nonnegativity. Because g is the sum of a signomial with all nonnegative coefficients (a *posynomial*) and a constant, it is convex by construction, its X -nonnegativity can be decided by applying the principle of strong duality in convex optimization. The outcome of this duality argument is that f is X -nonnegative if and only if there exists a dual variable $\nu = (\nu_\alpha)_{\alpha \in \mathcal{A}}$ that satisfies a certain relative entropy inequality in ν , c , and the support function of X (see Proposition 2.3 for a precise statement). Thus, the X -nonnegativity of f can be decided by convex *relative entropy programming*. The X -nonnegative signomials with at most one negative coefficient are called X -AGE, and the signomials which decompose into a sum of such functions are called X -SAGE. The recognition problem for X -SAGE signomials can likewise be decided by relative entropy programming.

The purpose of this article is to undertake the first structural analysis of the cones of X -SAGE signomials on exponents \mathcal{A} , which we henceforth denote by $C_X(\mathcal{A})$. At the outset of this research, our goals were to find counterparts to the many convex-combinatorial properties known for the unconstrained case $C_{\mathbb{R}^n}(\mathcal{A})$ [12, 17, 23], and to understand conditional SAGE relative to techniques such as nonnegative circuit polynomials [14, 32, 36]. Towards this end we have introduced an analysis tool of *sublinear circuits* which we call *the X -circuits of \mathcal{A}* . Our definition of these X -circuits (see Section 3) centers on a local, orthant-wise, strict-sublinearity condition for the support function of X composed with \mathcal{A} . This construction ensures that the special case of \mathbb{R}^n -circuits reduces to the simplicial circuits of the affine-linear matroid induced by \mathcal{A} .

We demonstrate that analysis by X -circuits is extremely effective in describing many structural aspects of X -SAGE cones. Our techniques are sufficiently robust that one can prove nearly every

result in this manuscript assuming nothing of X beyond convexity. Some special treatment is given to the case when X is polyhedral, as this reveals some striking interactions between discrete, convex, and so-called *geometrically convex* or *multiplicatively convex* geometry (see Section 5). In a broader sense, a selection of our results have consequences for numerical optimization, such as basis identification in optimization with SAGE certificates, and a procedure to simplify certain systems of power cone inequalities on the nonnegative orthant.

1.1. Main contributions. We begin by introducing some limited notation. The vector space of real $|\mathcal{A}|$ -tuples indexed by $\alpha \in \mathcal{A}$ is denoted by $\mathbb{R}^{\mathcal{A}}$. The support function of a convex set X , denoted by σ_X , is the convex function defined by $\sigma_X(y) = \sup\{y^T x : x \in X\}$. We regard the exponent set $\mathcal{A} \subset \mathbb{R}^n$ as a linear operator from $\mathbb{R}^{\mathcal{A}}$ to \mathbb{R}^n by $\mathcal{A}\nu = \sum_{\alpha \in \mathcal{A}} \alpha \nu_\alpha$; the corresponding adjoint is denoted \mathcal{A}^T . For concreteness, one might think of \mathcal{A} as a matrix with columns given by the exponents α . We continue to use $C_X(\mathcal{A})$ to denote the cone of X -SAGE signomials on \mathcal{A} . For each $\beta \in \mathcal{A}$, we denote the corresponding cone of X -AGE functions by

$$C_X(\mathcal{A}, \beta) = \left\{ f : f = \sum_{\alpha \in \mathcal{A}} c_\alpha e^\alpha \text{ is } X\text{-nonnegative, } c_{\setminus\beta} \geq \mathbf{0} \right\} \quad (2)$$

where $c_{\setminus\beta}$ denotes the vector in $\mathbb{R}^{\mathcal{A} \setminus \beta}$ formed by deleting c_β from c .

The basic tools for our analysis are *the X -circuits of \mathcal{A}* (routinely abbreviated to *X -circuits*). We formulate the X -circuits of \mathcal{A} as nonzero vectors $\nu^* \in \mathbb{R}^{\mathcal{A}}$ at which the augmented support function $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ exhibits a strict sublinearity condition (see Definition 3.1). We characterize X -circuits as generators of suitable convex cones in $\mathbb{R}^{\mathcal{A}} \times \mathbb{R}$ and usually focus on *normalized X -circuits* $\lambda \in \mathbb{R}^{\mathcal{A}}$, for which the nonnegative entries sum to unity. Theorem 3.7 shows that in the polyhedral case, X -circuits are exactly the generators of all one-dimensional elements of a suitable polyhedral fan. A key consequence of Theorem 3.7 is that when X is a polyhedron, there are only finitely many normalized X -circuits.

Section 4 uses the machinery of X -circuits to understand X -AGE cones. First, we show that if a signomial generates an extreme ray of $C_X(\mathcal{A}, \beta)$, then the dual variable ν which certifies its required relative entropy inequality must be an X -circuit (Theorem 4.2). Normalized X -circuits λ are then associated to cones of *λ -witnessed AGE functions* $C_X(\mathcal{A}, \lambda)$.¹ The functions in $C_X(\mathcal{A}, \lambda)$ are X -nonnegative signomials admitting a nonnegativity certificate based on a damped power cone inequality in weights λ . Theorem 4.4 shows that every X -SAGE function can be written as a sum of λ -witnessed AGE functions for X -circuits λ . In proving this, we formalize the connection between conditional SAGE and prior works for global nonnegativity [14, 32, 36]. Theorem 4.4 also motivates a basis identification technique where an approximate relative entropy certificate of $f \in C_X(\mathcal{A})$ may be refined by power cone programming. Combining Theorems 3.7 and 4.4 yields a corollary that when X is a polyhedron, cones of X -SAGE signomials are (in principle) power cone representable; this generalizes results by several authors in the unconstrained case [2, 26, 33, 41].

¹When parsing $C_X(\mathcal{A}, \beta)$ and $C_X(\mathcal{A}, \lambda)$, the reader should note that β and λ live in different spaces.

Section 5 undertakes a thorough analysis of $C_X(\mathcal{A})$. We begin by associating X -circuits λ with affine functions $\phi_\lambda : \mathbb{R}^A \rightarrow \mathbb{R}$ given by $\phi_\lambda(y) = \sum_{\alpha \in \mathcal{A}} y_\alpha \lambda_\alpha + \sigma_X(-\mathcal{A}\lambda)$. We define the *circuit-generated cone* $G_X(\mathcal{A})$ as the smallest convex cone containing these functions and the constant function $y \mapsto 1$. Upon embedding the affine functions on \mathbb{R}^A into $\mathbb{R}^A \times \mathbb{R}$, Theorem 5.4 provides the following identity between the dual SAGE cone $C_X(\mathcal{A})^*$ and the dual circuit-generated cone $G_X(\mathcal{A})^*$

$$C_X(\mathcal{A})^* = \text{cl}\{\exp y : (y, 1) \in G_X(\mathcal{A})^*\}.$$

Qualitatively, Theorem 5.4 says $C_X(\mathcal{A})^*$ is not only convex in the classical sense, but also convex under a logarithmic transformation $S \mapsto \log S = \{y : \exp y \in S\}$. The property of a set being convex under this logarithmic transformation is known by various names, including *log convexity* [1], *geometric convexity* [15, 30], or *multiplicative convexity* [28]. This property has previously been considered in the literature on ordinary SAGE certificates (i.e., SAGE certificates for the special case $X = \mathbb{R}^n$) [17, 23], but never in such a systematic way as in our analysis. For example, in view of Theorem 5.4 it becomes natural to consider $\Lambda_X^*(\mathcal{A})$ – the *reduced X -circuits of \mathcal{A}* – as the normalized circuits λ for which ϕ_λ generates an extreme ray of the circuit-generated cone. The property of a circuit being “reduced” in this sense is highly restrictive, and yet (by Theorem 5.5) we can construct $C_X(\mathcal{A})$ using only λ -witnessed AGE cones as λ runs over $\Lambda_X^*(\mathcal{A})$. Finally, through a technical lemma (5.14), we show how separating hyperplanes in the space of the dual circuit-generated cone may be mapped to separating hyperplanes in the exponentiated space of the dual SAGE cone. This lemma has general applications in simplifying systems of certain power cone constraints on the nonnegative orthant; in our context, it serves as the basis for Theorem 5.6, paraphrased below.

If X is a polyhedron and $C_X(\mathcal{A})$ consists of more than just posynomials, then

$$C_X(\mathcal{A}) = \sum_{\lambda \in \Lambda_X^*(\mathcal{A})} C_X(\mathcal{A}, \lambda).$$

Moreover, there is no subset $\Lambda \subsetneq \Lambda_X^(\mathcal{A})$ for which $C_X(\mathcal{A}) = \sum_{\lambda \in \Lambda} C_X(\mathcal{A}, \lambda)$.*

Theorem 5.6 provides the most efficient possible description of $C_X(\mathcal{A})$ in terms of power cone inequalities. Its computational implications are addressed briefly in Section 7.

Throughout the article we illustrate key concepts with the half-line $X = [0, \infty)$. Specifically, Example 3.8 addresses the $[0, \infty)$ -circuits of a generic point set $\mathcal{A} \subset \mathbb{R}$, and Example 5.7 covers the corresponding *reduced* $[0, \infty)$ -circuits. This culminates with a complete characterization of the extreme rays of $C_X(\mathcal{A})$ for $X = [0, \infty)$ and $\mathcal{A} \subset \mathbb{R}$ (Proposition 6.1).

1.2. Related work. Let us begin by introducing some basic concepts from discrete geometry. The *circuits* of the affine-linear matroid induced by \mathcal{A} are the nonzero vectors $\nu^* \in \ker \mathcal{A} \subset \mathbb{R}^A$ whose entries sum to zero, and whose supports are inclusion minimal among all vectors in $\ker \mathcal{A}$ that sum to zero. In the SAGE literature one is interested in *simplicial circuits*. These are the circuits ν^* that, upon scaling by a suitable constant, have exactly one negative component. The name *simplicial* is used here because the convex hull of the support $\text{supp } \nu^* := \{\alpha : \nu_\alpha^* \neq 0\}$ forms a simplex (possibly of low dimension); exactly one element in $\text{supp } \nu^*$ is contained in the relative interior of this simplex. These simplicial circuits are uniquely determined (up to

scaling) by their supports. It is therefore common to call a subset $A \subset \mathcal{A}$ a simplicial circuit if its convex hull forms a simplex and has a relative interior containing exactly one element of A .

To situate conditional SAGE in the literature one should look to the close relatives of ordinary SAGE: the *agiforms* of Reznick [36], the *monomial dominating posynomials* of Pantea, Koepl and Craciun [32], and the *Sums-of-Nonnegative-Circuit* (SONC) polynomials of Ilman and de Wolff [14]. The latter two works determined necessary and sufficient conditions for \mathbb{R}_+^n and \mathbb{R}^n -nonnegativity of polynomials supported on a simplicial circuit, based on power cone inequalities in the polynomial's coefficients and circuit vector. In our context, key developments in this area include Wang's discovery of conditions under which a SONC decomposition exists for a given polynomial [40], and Murray, Chandrasekaran, and Wierman's proof that the cone of SONC polynomials can be represented by a projection of a cone of SAGE signomials [23, §5]. From these results it is now understood that SONC and ordinary SAGE are equivalent to one another for purposes of certain structural analyses. Our results show that the "circuit number" approach of SONC does not generalize to the X -nonnegativity problem in the same manner as SAGE. However, it is possible to describe conditional SAGE in a way which is aesthetically similar to SONC via our λ -witnessed AGE cones.

To appreciate the structural results proven for $C_X(\mathcal{A})$ in this work, it is useful to mention some analogous results proven in the case $X = \mathbb{R}^n$. As a signomial generalization of an earlier result by Reznick [36], Murray, Chandrasekaran, and Wierman have shown that every signomial which generates an extreme ray of $C_{\mathbb{R}^n}(\mathcal{A})$ is supported on either a singleton or a simplicial circuit [23]. Curiously, a given signomial f can be extremal in $C_{\mathbb{R}^n}(\mathcal{A})$ for \mathcal{A} as the support of f , and yet nonextremal in $C_{\mathbb{R}^n}(\mathcal{A}')$ for $\mathcal{A}' \supsetneq \mathcal{A}$. To account for this, Katthän, Naumann, and Theobald introduced the concept of a *reduced circuit*, which they used to obtain a complete characterization of the extreme rays of $C_{\mathbb{R}^n}(\mathcal{A})$ [17]. Subsequently, Forsgård and de Wolff employed regular subdivisions, A -discriminants and tropical geometry to study how circuits affect the algebraic boundary of the signomial SAGE cone [12]. Our results include direct extensions of the above results by Murray et al. and Katthän et al. to the case of $X \subsetneq \mathbb{R}^n$. For Forsgård and de Wolff's work, our *circuit-generated cone* generalizes their *Reznick cone*.

Now we turn to how SAGE can be used for optimization. Given a signomial objective f and a convex feasible set X , we have $\sup\{\gamma \in \mathbb{R} : f - \gamma \text{ is } X\text{-SAGE}\} \leq \inf_{x \in X} f(x)$. This procedure has been extended to a convex relaxation hierarchy for which A. Wang et al. have proven a completeness result [39] (see also [6]). Very recently, additional SAGE-based hierarchies have been developed to approach a signomial's minimum from both above and below, including in the presence of nonconvex constraints [9]. Such techniques can be implemented using the `sageopt` python package and a reliable exponential cone solver such as MOSEK [5, 21].

On the polynomial optimization side, Karaca et al. developed a combined SAGE and Sums-of-Squares approach to optimization over (subsets of) the nonnegative orthant [16]. By consideration to the close SAGE-SONC relationship, one finds connections to works of Dressler et al. on polynomial optimization with SONC [7, 8]. As an alternative to SONC, one may work directly with a notion of *SAGE polynomials* [23, §5.1]. The concept of SAGE polynomials is important because the corresponding nonnegativity certificates can be computed efficiently, and because they are transparently generalized to *X -SAGE polynomials* [24, §4]. Our signomial results may

be applied to conditional SAGE polynomials, however care must be taken in mapping between the two types of functions; see for example [24, Theorems 1 and 2].

1.3. Some definitions and conventions. Our terminology and notation for convex analysis is generally chosen to match that of Rockafellar [37]. Here we define terms and notation which are less commonly used or which differ from those of [37]; additional standard definitions are reproduced in the appendix. We abbreviate the line segment connecting x and y in \mathbb{R}^n by $[x, y] = \{\lambda x + (1 - \lambda)y : 0 \leq \lambda \leq 1\}$. A convex cone $K \subset \mathbb{R}^n$ is *pointed* if it contains no lines. A vector v in a convex cone K is called an *edge generator* if $\{\lambda v : \lambda \geq 0\}$ is an extreme ray of K . The *polar* of a convex cone K is $K^\circ = -K^*$, where K^* is the dual cone to K . The *induced cone* of a convex set $S \subset \mathbb{R}^n$ is $\text{indco}(S) := \text{cl}\{(s, \mu) : \mu > 0, s/\mu \in S\} \subset \mathbb{R}^{n+1}$, and the *recession cone* is $\text{rec}(S) := \{t : \exists s \in S \text{ such that } s + \lambda t \in S \forall \lambda \geq 0\}$.

All logarithms are base- e , where e is Euler’s number. We extend the scalar exponential function “exp” to real vectors in an elementwise fashion. The zero vector and vector of all ones (in appropriate spaces) are denoted $\mathbf{0}$ and $\mathbf{1}$ respectively. The standard basis for \mathbb{R}^A is denoted $\{\delta_\alpha\}_{\alpha \in A}$, and the support of a vector $c \in \mathbb{R}^A$ is $\text{supp } c = \{\alpha : c_\alpha \neq 0\}$.

ACKNOWLEDGEMENTS

This work would not have been possible without an invitation from Bernd Sturmfels for R.M. to visit The Max Planck Institute for Mathematics in the Sciences (Leipzig, Germany) in late 2019. R.M. was supported by an NSF Graduate Research Fellowship, and T.T. was supported by DFG grant TH 1333/7-1. We thank the anonymous referees for their constructive feedback.

2. PRELIMINARIES

Throughout this article, $X \subset \mathbb{R}^n$ is closed, convex, and nonempty, and the set $\mathcal{A} \subset \mathbb{R}^n$ is nonempty and finite. We only consider data (\mathcal{A}, X) where the functions $\{e^\alpha\}_{\alpha \in \mathcal{A}}$ are linearly independent on X . The purpose of this linear independence assumption is to ensure the X -nonnegativity cone does not contain a lineality space; equivalently, the assumption ensures the moment cone $\text{co}\{\exp(\mathcal{A}^T x) \in \mathbb{R}^{\mathcal{A}} : x \in X\}$ is full-dimensional.

Definition 2.1. The X -SAGE cone with respect to the support \mathcal{A} is the Minkowski sum

$$C_X(\mathcal{A}) = \sum_{\beta \in \mathcal{A}} C_X(\mathcal{A}, \beta),$$

where $C_X(\mathcal{A}, \beta)$ are the X -AGE cones defined in (2).

Remark 2.2. In this definition, all the signomials in the decomposition of the right hand side are also restricted to the support \mathcal{A} . This is no loss of generality, since any signomial f on \mathcal{A} , which is contained in $\sum_{\beta \in \mathcal{A}'} C_X(\mathcal{A}', \beta)$ for some superset \mathcal{A}' of \mathcal{A} , is also contained in $\sum_{\beta \in \mathcal{A}} C_X(\mathcal{A}, \beta)$, see [24, Corollary 1].

By adopting Definition 2.1, it is clear that the problem of representing $C_X(\mathcal{A})$ reduces to the problem of representing the cones $C_X(\mathcal{A}, \beta)$. To state the representation of these cones we use the *relative entropy function*

$$D(\nu, c) = \sum_{\alpha \in \mathcal{A}} \nu_\alpha \log \left(\frac{\nu_\alpha}{c_\alpha} \right).$$

We use standard conventions where relative entropy is continuously extended to $\mathbb{R}_+^{\mathcal{A}} \times \mathbb{R}_+^{\mathcal{A}}$, and define $D(\nu, c) = \infty$ if either ν or c has a negative component.

Proposition 2.3 (Theorem 1 of [24]). *A signomial $f = \sum_{\alpha \in \mathcal{A}} c_\alpha e^\alpha$ belongs to $C_X(\mathcal{A}, \beta)$ if and only if there exists a vector $\nu \in \mathbb{R}^{\mathcal{A}}$ that satisfies*

$$\mathbf{1}^T \nu = 0 \quad \text{and} \quad \sigma_X(-\mathcal{A}\nu) + D(\nu_{\setminus \beta}, e_{c_{\setminus \beta}}) \leq c_\beta, \quad (3)$$

where again, $\sigma_X(y) = \sup\{y^T x : x \in X\}$ for $y \in \mathbb{R}^n$. Such a vector ν is called a *relative entropy certificate* for f .

Proposition 2.3 is important for computational optimization. For example, if X is the unit ball in the Euclidean norm, then $\sigma_X(-\mathcal{A}\nu) = \|\mathcal{A}\nu\|_2$, and so (3) becomes a mixed relative-entropy and second-order-cone inequality. More generally, the formulation is tractable whenever we can efficiently represent the epigraph of the support function of X .

Since the relative entropy condition in Proposition 2.3 is essential for our treatment, we outline its proof. Adopt I_X as the indicator function of X , with $I_X(x) = 0$ for $x \in X$ and $I_X(x) = \infty$ otherwise. Given $f = \sum_{\alpha \in \mathcal{A}} c_\alpha e^\alpha$ with $c_{\setminus \beta} \geq \mathbf{0}$, the primal formulation for X -nonnegativity of f is

$$\inf_{\substack{x \in \mathbb{R}^n \\ t \in \mathbb{R}^{\mathcal{A}}}} \left\{ I_X(x) + \sum_{\alpha \in \mathcal{A} \setminus \beta} c_\alpha \exp t_\alpha : t_\alpha = (\alpha - \beta)^T x \quad \forall \alpha \in \mathcal{A} \right\} \geq -c_\beta. \quad (4)$$

The formulation (3) is simply the dual to (4) using the machinery of convex conjugate functions. In particular, the relative entropy certificate ν in (3) is the dual variable to the equality constraints in (4).

The larger goal of this article is to reveal additional structure in the X -SAGE cones $C_X(\mathcal{A})$ that is not immediately apparent from Proposition 2.3. From the case $X = \mathbb{R}^n$, the additional structure concerned the supports of signomials that generate extreme rays of $C_{\mathbb{R}^n}(\mathcal{A}, \beta)$ or $C_{\mathbb{R}^n}(\mathcal{A})$. In this context it is standard to use the term *simplicial circuit* in the sense of subsets $A \subset \mathcal{A}$. Specifically, $A \subset \mathcal{A}$ is a simplicial circuit if it is a minimal affinely dependent set and $\text{conv } A$ has $|A| - 1$ extreme points. This definition of circuits in terms of these subsets $A \subset \mathcal{A}$ is equivalent to the definition involving numeric vectors $\nu^* \in \mathbb{R}^{\mathcal{A}}$; see [12].

Proposition 2.4 (Theorem 5 of [23]). *Let $\beta \in \mathcal{A}$. A signomial $f = \sum_{\alpha \in \mathcal{A}} c_\alpha e^\alpha$ belongs to $C_{\mathbb{R}^n}(\mathcal{A}, \beta)$ if and only if it can be written as a finite sum $f = \sum_{i=1}^k f^{(i)}$ of signomials*

$$f^{(i)} = \sum_{\alpha \in \mathcal{A}} c_\alpha^{(i)} e^\alpha \in C_{\mathbb{R}^n}(\mathcal{A}, \beta), \quad 1 \leq i \leq k,$$

such that the supports $\{\alpha \in \mathcal{A} : c_\alpha^{(i)} \neq 0\}$ are either singletons or simplicial circuits.

Of course, in view of Definition 2.1, Proposition 2.4 tells us every $f \in C_{\mathbb{R}^n}(\mathcal{A})$ similarly decomposes into AGE functions supported on singletons and simplicial circuits.

Revealing the full structure of conditional SAGE cones requires consideration to more than just a signomial's support. Therefore, thinking in terms of affine-linear circuits as subsets $A \subset \mathcal{A}$ will not suit our purposes. The following definition codifies our convention of considering affine-linear circuits as numeric vectors.

Definition 2.5. A nonzero vector $\nu^* \in \{\nu \in \mathbb{R}^{\mathcal{A}} : \mathbf{1}^T \nu = 0\}$ in the kernel of the linear operator $\nu \mapsto \mathcal{A}\nu = \sum_{\alpha \in \mathcal{A}} \alpha \nu_{\alpha}$ is called an \mathbb{R}^n -circuit if it is minimally supported and has exactly one negative component.

It is possible that a given \mathcal{A} has no \mathbb{R}^n -circuits, but then every $\alpha \in \mathcal{A}$ would be an extreme point of $\text{conv } \mathcal{A}$. This is a degenerate case that results in $C_{\mathbb{R}^n}(\mathcal{A})$ containing only posynomials, but we still give consideration to this possibility throughout the article. In the language of Definition 2.5, we combine Propositions 2.3 and 2.4 to obtain the following formulation.

Proposition 2.6 (Theorem 4.4 of [12]). *Let $\beta \in \mathcal{A}$. A signomial $f = \sum_{\alpha \in \mathcal{A}} c_{\alpha} e^{\alpha}$ belongs to $C_{\mathbb{R}^n}(\mathcal{A}, \beta)$ if and only if there exist $k \geq 0$ and signomials $f^{(i)} = \sum_{\alpha \in \mathcal{A}} c_{\alpha}^{(i)} e^{\alpha} \in C_{\mathbb{R}}(\mathcal{A}, \beta)$, $1 \leq i \leq k$, with $f = \sum_{i=1}^k f^{(i)}$ and such that for any signomial $f^{(i)}$ which is not supported on a singleton, there exists an \mathbb{R}^n -circuit $\nu^{(i)} \in \mathbb{R}^{\mathcal{A}}$ with $D(\nu_{\setminus \beta}^{(i)}, e c_{\setminus \beta}^{(i)}) \leq c_{\beta}^{(i)}$.*

3. SUBLINEAR CIRCUITS INDUCED BY A POINT SET

We begin this section with a functional analytic definition for the X -circuits of a point set \mathcal{A} , generalizing \mathbb{R}^n -circuits to a constrained setting. After revealing various elementary properties and discussing some examples, we characterize X -circuits in more geometric terms in Theorems 3.6 and 3.7. In particular the latter theorem interprets X -circuits in terms of normal fans when X is a polyhedron. In Example 3.8, we determine the $[0, \infty)$ -circuits of a univariate support set $\mathcal{A} \subset \mathbb{R}$; the example is developed further in Section 5 and culminates in a theorem completely characterizing the extreme rays of the resulting X -SAGE cone $C_{[0, \infty)}(\mathcal{A})$ in Section 6.

The derivations in this section are purely combinatorial and convex-geometric, and make no mention of signomials. However, the definition of X -circuits is ultimately chosen to prepare for studying X -SAGE cones, and in particular it relates to distinguished vectors $\nu \in \mathbb{R}^{\mathcal{A}}$ that might satisfy (3) for certain $c \in \mathbb{R}^{\mathcal{A}}$. Note that (3) has an implicit constraint $\nu_{\setminus \beta} \geq \mathbf{0}$ arising from our extended-real-valued definition of relative entropy. To avoid dependence on relative entropy in this section, we frame our discussion of X -circuits in terms of cones

$$N_{\beta} = \{\nu \in \mathbb{R}^{\mathcal{A}} : \nu_{\setminus \beta} \geq \mathbf{0}, \mathbf{1}^T \nu = 0\} \quad (5)$$

for vectors $\beta \in \mathcal{A}$.

Definition 3.1. A vector $\nu^* \in N_{\beta}$ is an X -circuit of \mathcal{A} (or simply, an X -circuit) if (1) it is nonzero, (2) $\sigma_X(-\mathcal{A}\nu^*) < \infty$, and (3) it cannot be written as a convex combination of two non-proportional $\nu^{(1)}, \nu^{(2)} \in N_{\beta}$, for which $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ is linear on $[\nu^{(1)}, \nu^{(2)}]$.

The third condition is equivalent to strict sublinearity of $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ on any line segment in N_{β} that contains ν^* , except for the trivial line segments which generate a single ray. The

central importance of the sublinearity condition leads us to refer to X -circuits also as *sublinear circuits*; the latter term is helpful in remembering the definition early in our development.

Remark 3.2. In the special case $X = \mathbb{R}^n$, condition (2) simplifies to $\mathcal{A}\nu = \mathbf{0}$. In conjunction with the definition of N_β , this shows that the special case $X = \mathbb{R}^n$ of Definition 3.1 matches exactly with Definition 2.5 of \mathbb{R}^n -circuits.

Conceptually, Definition 3.1 indicates that X -circuits are essential in capturing the behavior of the augmented support function $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ on the given N_β . While developing this concept formally it is convenient for us to enumerate $\nu^+ := \{\alpha : \nu_\alpha > 0\}$, and to identify the unique index $\nu^- := \beta \in \mathcal{A}$ where $\nu_\beta < 0$. Note that positive homogeneity of the support function tells us that the property of being a sublinear circuit is invariant under scaling by positive constants. A sublinear circuit is *normalized* if its unique negative term ν_β has $\nu_\beta = -1$, in which case we usually denote it by the symbol λ rather than ν . We can normalize a given sublinear circuit by taking the ratio with its infinity norm $\lambda = \nu/\|\nu\|_\infty$, because $\|\nu\|_\infty = |\nu_\beta|$ for all vectors $\nu \in N_\beta$.

Example 3.3. (*The conic case.*) It is straightforward to determine which $\nu \in N_\beta$ are X -circuits of \mathcal{A} when X is a cone. In such a setting, the support function of X can only take on the values zero and positive infinity. Hence, $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ is trivially linear over all of $V_\beta := \{\nu \in N_\beta : \sigma_X(-\mathcal{A}\nu) < \infty\}$. Notice that V_β is a cone and that $\sigma_X(-\mathcal{A}\nu) = 0$ may be reformulated as $\nu \in (\mathcal{A}^T X)^*$. Standard conic duality calculations (see Proposition 8.3) show that $(\mathcal{A}^T X)^* = \ker \mathcal{A} + \mathcal{A}^\dagger X^*$, where \mathcal{A}^\dagger denotes the Moore-Penrose pseudo-inverse of \mathcal{A} . Thus

$$V_\beta = (\ker \mathcal{A} + \mathcal{A}^\dagger X^*) \cap N_\beta$$

and the X -circuits $\nu \in N_\beta$ are precisely the edge generators of V_β .

Regarding again the special case $X = \mathbb{R}^n$ from this conic perspective, we have $X^* = \{\mathbf{0}\}$, so $\mathcal{A}^\dagger X^* = \{\mathbf{0}\}$, and $\ker \mathcal{A} + \mathcal{A}^\dagger X^* = \ker \mathcal{A}$, which implies $V_\beta = \ker \mathcal{A} \cap N_\beta$. It is easily shown that edge generators of $\ker \mathcal{A} \cap N_\beta$ are precisely those $\nu \in \ker \mathcal{A} \cap N_\beta \setminus \{\mathbf{0}\}$ for which $\nu^+ = \{\alpha : \nu_\alpha > 0\}$ are affinely independent, which recovers the matroid-theoretic notion of affine-linear simplicial circuits from the point of view of subsets $A \subset \mathcal{A}$.

The following proposition shows that the affine-independence property is a necessary condition for all sublinear circuits. The proposition provides insight because it shows an X -circuit ν with $X \subset \mathbb{R}^n$ is restricted to $|\text{supp } \nu| \leq n + 2$.

Proposition 3.4. *If $\nu^* \in N_\beta$ is an X -circuit, then $(\nu^*)^+ = \text{supp } \nu^* \setminus \beta$ is affinely independent.*

Proof. From a fixed $\nu^* \in N_\beta$ construct $z = -\mathcal{A}\nu^*$ and $U = \{\nu \in N_\beta : -\mathcal{A}\nu = z, \nu_\beta = \nu_\beta^*\}$. The function $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ is a constant and equal to $\sigma_X(z)$ on U , and so in order for ν^* to be an X -circuit, it must be a vertex of the polytope U . The set U is in 1-to-1 correspondence with $W = \{w \in \mathbb{R}_+^{\mathcal{A} \setminus \beta} : \sum_{\alpha \in \mathcal{A} \setminus \beta} (\beta - \alpha)w_\alpha = z, \mathbf{1}^T w = -\nu_\beta^*\}$ by identifying $w = \nu_{\setminus \beta}$. In matrix notation, we can write $W = \{w \in \mathbb{R}_+^{\mathcal{A} \setminus \beta} : Mw = (z, -\nu_\beta^*)\}$ by forming the matrix M with columns $\{(\beta - \alpha, 1)\}_{\alpha \in \mathcal{A} \setminus \beta}$ indexed by $\alpha \in \mathcal{A} \setminus \beta$.

Basic polyhedral geometry tells us that all vertices w^* of W use an affinely independent set of columns from M . Furthermore, a given set of columns from M is affinely independent if and

only if the corresponding indices of the columns (as vectors $\alpha \in \mathcal{A} \setminus \beta$) are affinely independent. Since the correspondence between $\nu \in U$ and $w \in W$ preserves extremality, the vertices of U have affinely independent positive support ν^+ . \square

The converse of Proposition 3.4 is not true. This is to say: not every vector $\nu \in N_\beta$ with affinely independent ν^+ is an X -circuit.

Example 3.5. Let $\mathcal{A} \subset \mathbb{R}^2$ contain $\alpha_1 = (0, 0)$, $\alpha_2 = (1, 0)$, and $\alpha_3 = (0, 1)$, and consider $X = \{x \in \mathbb{R}^2 : x \geq u\}$ for some fixed point $u \in \mathbb{R}^2$. The vector $\nu^* = (-2, 1, 1)$ has $(\nu^*)^- = \alpha_1 = (0, 0)$, and $(\nu^*)^+ = \{\alpha_2, \alpha_3\} = \{(1, 0), (0, 1)\}$ is affinely independent. Considering $\nu^{(1)} = (-2, 2, 0)$ and $\nu^{(2)} = (-2, 0, 2)$, we have $\nu^* = \frac{1}{2}(\nu^{(1)} + \nu^{(2)}) \in \text{ri} L$ for $L := [\nu^{(1)}, \nu^{(2)}]$. Moreover, the mapping $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ is linear on L , because for any $\mu_1, \mu_2 \geq 0$ with $\mu_1 + \mu_2 = 1$ we have

$$\begin{aligned} \sigma_X(\mathcal{A}(-\mu_1\nu^{(1)} - \mu_2\nu^{(2)})) &= \sigma_X((-2\mu_1, -2\mu_2)) = -2\mu_1u_1 - 2\mu_2u_2 \\ &= \sigma_X((-2\mu_1, 0)) + \sigma_X((0, -2\mu_2)). \end{aligned}$$

The last equality is true since $(1, 1)$ maximizes both the objective functions $x \mapsto (-2\mu_1, 0)^T x$ and $x \mapsto (0, -2\mu_2)^T x$ on X .

With the basic exercise of Example 3.5 complete, we turn to characterizing sublinear circuits in full generality.

Theorem 3.6. Fix $\beta \in \mathcal{A}$. The convex cone generated by

$$T = \{(\nu, \sigma_X(-\mathcal{A}\nu)) : \nu \in N_\beta, \sigma_X(-\mathcal{A}\nu) < \infty\}$$

is pointed (i.e., it contains no lines) and closed. A vector $\nu^* \in N_\beta$ is an X -circuit of \mathcal{A} if and only if $(\nu^*, \sigma_X(-\mathcal{A}\nu^*))$ is an edge generator for $\text{co}T$.

Proof. Let Q denote the closed convex set $Q = \{\nu : \nu \in N_\beta, \sigma_X(-\mathcal{A}\nu) < \infty\}$. The claim of the theorem is trivially true if $Q = \{\mathbf{0}\}$, in which case there are no X -circuits $\nu \in N_\beta$ and $\text{co}T = \{(\mathbf{0}, 0)\}$ has no extreme rays. We therefore assume for the duration of the proof that Q contains a nonzero vector.

We turn to showing $\text{co}T$ is closed and pointed, particularly beginning with pointedness. For this, observe $\text{co}T \subset N_\beta \times \mathbb{R}$. Since N_β contains no lines, there are no lines in $\text{co}T$ of the form (ν, τ) with $\nu \neq \mathbf{0}$. Meanwhile, we know that the line spanned by $(\mathbf{0}, 1)$ cannot be contained in $\text{co}T$, since $\sigma_X(-\mathcal{A}\mathbf{0}) = 0$. Now we turn to closedness of $\text{co}T$. Since Q is contained within N_β , we may normalize Q against $\{\nu : \nu_\beta = -1\}$: $Q = \text{co}Q_1$ for the nonempty compact convex set $Q_1 := \{\lambda : \lambda \in Q, \lambda_\beta = -1\}$. From Q_1 we construct $T_1 = \{(\lambda, \sigma_X(-\mathcal{A}\lambda)) : \lambda \in Q_1\}$. The set T_1 inherits compactness from Q_1 (by continuity of $\lambda \mapsto \sigma_X(-\mathcal{A}\lambda)$), and the convex hull $T_2 = \text{conv}T_1$ inherits compactness from T_1 (as the convex hull of a compact set is compact). It is evident that T_2 does not contain the zero vector, and so by [37, Corollary 9.6.1] we have that $\text{co}T_2$ is closed. We finish this phase of the proof by identifying $\text{co}T = \text{co}T_2$.

At this point we have that $\text{co}T$ is the convex hull of its extreme rays; it remains to determine the nature of these extreme rays. Since T is a generating set for $\text{co}T$ and contains only vectors of the form $(\nu, \sigma_X(-\mathcal{A}\nu))$, every edge generator of $\text{co}T$ is given by a nonzero vector $(\nu^*, \sigma_X(-\mathcal{A}\nu^*))$

for appropriate ν^* . It is clear that ν^* must be an X -circuit in order for $(\nu^*, \sigma_X(-\mathcal{A}\nu^*))$ to be an edge generator of $\text{co}T$. The harder direction is to show that ν^* being an X -circuit is *sufficient* for $(\nu^*, \sigma_X(-\mathcal{A}\nu^*))$ to be an edge generator for $\text{co}T$.

To handle this direction, begin by defining an affinely independent set $\mathcal{V} = \{\nu^{(i)}\}_{i=1}^{\ell}$ and a vector θ in the relative interior of $\Delta_{\ell} := \{z \in \mathbb{R}_+^{\ell} : \mathbf{1}^T z = 1\}$, where $\nu^* = \sum_{i=1}^{\ell} \theta_i \nu^{(i)}$ and

$$\sigma_X(-\mathcal{A}\nu^*) = \sum_{i=1}^{\ell} \theta_i \sigma_X(-\mathcal{A}\nu^{(i)}).$$

We claim that $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ is linear on the entirety of $\text{conv } \mathcal{V}$. To see why, note that the assumption on ν^* relative to \mathcal{V} means the elements of $\Phi := \{(\nu^{(i)}, \sigma_X(-\mathcal{A}\nu^{(i)})) : i \in [\ell] \cup \{\star\}\}$ lie on a common hyperplane on the boundary of the epigraph $H = \{(\nu, t) : \sigma_X(-\mathcal{A}\nu) \leq t\}$. Since $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ is convex, H is a convex set, and there is some proper face F of H which contains Φ . It is evident that $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ is linear on the projection of that face $\hat{F} = \{\nu : \exists t \in \mathbb{R} (\nu, t) \in F\}$. Since $\text{conv } \mathcal{V} \subset \hat{F}$, this proves our claim regarding linearity of $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ on $\text{conv } \mathcal{V}$.

By the above argument: if ν^* is an X -circuit, then for every $\theta \in \text{ri } \Delta_{\ell}$ and affinely independent $\mathcal{V} = \{\nu^{(i)}\}_{i=1}^{\ell} \subset N_{\beta}$ with $\text{co } \mathcal{V} \neq \text{co}\{\nu^*\}$, we have

$$(\nu^*, \sigma_X(-\mathcal{A}\nu^*)) \neq \sum_{i=1}^{\ell} \theta_i (\nu^{(i)}, \sigma_X(-\mathcal{A}\nu^{(i)})).$$

From Carathéodory's Theorem, restricting to affinely independent $\mathcal{V} \subset T$ is sufficient to test extremality in $\text{co}T$. Therefore, every circuit $\nu^* \in N_{\beta}$ induces an edge generator for $\text{co}T$. \square

When considering the set “ T ” in Theorem 3.6, it is natural to expect that for polyhedral X there are only finitely many extreme rays in the cone $\text{co}T$, and hence only finitely many normalized X -circuits. The remainder of this section serves to prove this fact; here we use the concept of *normal fans* from polyhedral geometry. See, e.g., [43, Chapter 7] (for the bounded case of polytopes), [13, Section 5.4] or [38, Chapter 2]. For each face F of a polyhedron P , there is an associated *outer normal cone*

$$\mathbf{N}_P(F) = \{w : z^T w = \sigma_P(w) \ \forall z \in F\}.$$

Clearly, the support function of a polyhedron P is linear on every outer normal cone, and in particular the linear representation may be given by $\sigma_P(w) = z^T w$ for any $z \in F$. We obtain the *outer normal fan* of P by collecting all outer normal cones:

$$\mathcal{O}(P) = \{\mathbf{N}_P(F) : F \text{ is a face of } P\}$$

The support of $\mathcal{O}(P)$ is the polar $\text{rec}(P)^{\circ}$. The full-dimensional linearity domains of the support function are the outer normal cones of the vertices of P (see also [11, Section 1]).

Theorem 3.7. *If X is polyhedral, then $\nu \in N_{\beta} \setminus \{0\}$ is an X -circuit if and only if $\text{co}\{\nu\}$ is a ray in $\mathcal{O}(-\mathcal{A}^T X + N_{\beta}^{\circ})$. Consequently, a polyhedral set X has finitely many normalized circuits.*

Proof. Let $P = -\mathcal{A}^T X + N_{\beta}^{\circ}$. Using the characterization in [37, Theorem 14.2], the polar of its recession cone can be expressed as

$$(\text{rec } P)^{\circ} = \{\nu : \sigma_X(-\mathcal{A}\nu) < \infty\} \cap N_{\beta},$$

where we have also used the property $\sigma_X(-\mathcal{A}\nu) = \sup_{x \in X} (-\mathcal{A}\nu)^T x = \sup_{x \in -\mathcal{A}^T X} \nu^T x = \sigma_{-\mathcal{A}^T X}(\nu)$. In particular, this also gives $\sigma_X(-\mathcal{A}\nu) = \sigma_P(\nu)$. From P construct the outer normal fan $\mathcal{O} := \mathcal{O}(P)$. We claim that $\text{co}\{\nu\}$ is a ray in \mathcal{O} .

It is clear that if a cone $K \in \mathcal{O}$ is associated to a face F of P , then we may express $\sigma_P(\nu) = z^T \nu$ for any $z \in F$, and so $\sigma_P(\nu) \equiv \sigma_X(-\mathcal{A}\nu)$ is linear on K . Since the support of \mathcal{O} is $\text{rec}(P)^\circ$, the cones $K \in \mathcal{O}$ partition $(\text{rec } P)^\circ$, i.e.,

$$(\text{rec } P)^\circ = \bigcup_{K \in \mathcal{O}} \text{ri}(K),$$

and if K, K' are distinct elements in \mathcal{O} , then $\text{ri } K \cap \text{ri } K' = \emptyset$. Therefore, every $\nu \in N_\beta \setminus \{\mathbf{0}\}$ for which $\sigma_X(-\mathcal{A}\nu) < \infty$ is associated with a unique $K \in \mathcal{O}$, by way of $\nu \in \text{ri } K$.

Fix $\nu \in (\text{rec } P)^\circ$, and let K be the associated element of \mathcal{O} that contains ν in its relative interior. If K is of dimension greater than 1, ν can be expressed as a convex combination of non-proportional $\nu^{(1)}, \nu^{(2)} \in K$ – and clearly $\bar{\nu} \mapsto \sigma_X(-\mathcal{A}\bar{\nu}) \equiv \sigma_P(\bar{\nu})$ would be linear on the interval $[\nu^{(1)}, \nu^{(2)}]$. Thus for ν to be an X -circuit, it is necessary that K be of dimension 1. Since P is a polyhedron, \mathcal{O} is induced by finitely many faces. Thus there are finitely many $K \in \mathcal{O}$ with $\dim K = 1$ and in turn finitely many normalized X -circuits of \mathcal{A} .

Conversely, let $\nu^* \in N_\beta \setminus \{\mathbf{0}\}$ and $\text{co}\{\nu^*\}$ be a ray in \mathcal{O} . Since \mathcal{O} is supported on $\text{rec}(P)^\circ$, we have $\sigma_X(-\mathcal{A}\nu^*) = \sigma_P(\nu^*) < \infty$.

Let $\nu^{(1)}, \nu^{(2)} \in N_\beta$ be non-proportional and $\tau \in (0, 1)$ satisfy $\nu^* = \tau\nu^{(1)} + (1 - \tau)\nu^{(2)}$. If $\nu^{(1)}$ or $\nu^{(2)}$ is outside of $\text{rec}(P)^\circ$, say, $\nu^{(1)}$, then $\sigma_X(-\mathcal{A}\nu^{(1)}) = \infty$ and thus the mapping $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ cannot be linear on $[\nu^{(1)}, \nu^{(2)}]$. Hence, we can assume that $\nu^{(1)}, \nu^{(2)} \in \text{rec}(P)^\circ$.

We have to show that the mapping

$$g : [0, 1] \rightarrow \mathbb{R}, \quad \theta \mapsto \sigma_P(\theta\nu^{(1)} + (1 - \theta)\nu^{(2)})$$

is not linear.

Consider the restriction of the fan \mathcal{O} to the cone $C := \text{co}\{\nu^{(1)}, \nu^{(2)}\}$, that is, the collection of all the cones in $\{\mathcal{N}_P(F) \cap C : F \text{ is a face of } P\}$. This is a fan \mathcal{O}' supported on the two-dimensional cone $S := \text{rec}(P)^\circ \cap C$. On the set S , we consider the restricted mapping $(\sigma_P)|_S : S \rightarrow \mathbb{R}$, $w \mapsto \sigma_P(w)$. The linearity domains of $(\sigma_P)|_S$ are the two-dimensional cones in \mathcal{O}' . Since $\text{co}\{\nu^*\}$ is a ray in the fan \mathcal{O} and thus also in the fan \mathcal{O}' , the vectors $\nu^{(1)}$ and $\nu^{(2)}$ are contained in different two-dimensional cones of the fan \mathcal{O}' . Hence, the mapping g is not linear. Altogether, this shows that ν^* is an X -circuit. \square

Example 3.8. We consider as a running example the one-dimensional case of $X = [0, \infty)$ and $\mathcal{A} = \{\alpha_1, \dots, \alpha_m\} \subset \mathbb{R}$ where we can assume $\alpha_1 < \dots < \alpha_m$. In this running example we index by integers $i \in [m] := \{1, \dots, m\}$ rather than by elements $\alpha \in \mathcal{A}$. Therefore we identify $\mathbb{R}^{\mathcal{A}}$ with \mathbb{R}^m and use δ_i for the i^{th} unit vector in \mathbb{R}^m (for each $i \in [m]$). Under these conventions, \mathcal{A} is regarded as a row vector in $\mathbb{R}^{1 \times m}$ and $\mathcal{A}^T = (\alpha_1, \dots, \alpha_m)$ is a column vector in \mathbb{R}^m . We claim that the normalized X -circuits $\lambda \in \mathbb{R}^m$ are the vectors either of the form (1) $\lambda = \delta_k - \delta_j$ for $j < k$ or of the form (2)

$$\lambda = \left(\frac{\alpha_j - \alpha_i}{\alpha_k - \alpha_i} \right) \delta_k + \left(\frac{\alpha_k - \alpha_j}{\alpha_k - \alpha_i} \right) \delta_i - \delta_j \quad \text{for } i < j < k.$$

Note that vectors of type (2) satisfy $\mathcal{A}\lambda = 0$, and in fact are the unique such vectors that also satisfy $\text{supp } \lambda = \{i, j, k\}$, $\lambda_j = -1$, $\lambda_i, \lambda_k > 0$, $\mathbf{1}^T \lambda = 0$.

To derive this claim we consider for fixed $j \in [m]$ the polyhedron $P = -\mathcal{A}^T X + N_j^\circ$ from Theorem 3.7. It is evident that this polyhedron is a cone, that may be expressed as

$$P = \text{co}\{(-\alpha_1, \dots, -\alpha_m)\} + \mathbb{R} \cdot \mathbf{1} - \sum_{\ell \in [m] \setminus j} \text{co}\{\delta_\ell\}.$$

The rays of its normal fan are the extreme rays of its polar

$$P^\circ = (\text{rec } P)^\circ = \{\nu \in \mathbb{R}^m : (-\alpha_1, \dots, -\alpha_m)^T \nu \leq 0, \mathbf{1}^T \nu = 0, \nu_\ell \geq 0 \text{ for } \ell \in [m] \setminus j\}. \quad (6)$$

Note here that this gives us exactly the set “ Q ” from the proof of Theorem 3.6. This happens because X is conic and hence the support function $\sigma_X(-\mathcal{A}\nu)$ evaluates to zero for every X -circuit ν . By Proposition 3.4, each X -circuit in N_j has at most three non-vanishing components ν_i, ν_j, ν_k , and, moreover, it has $m - 2$ of the inequalities in (6) binding. If all those binding inequalities are of the form $\nu_\ell \geq 0$, then with $\sigma_X(-\mathcal{A}\nu) < \infty$, we obtain the normalized X -circuits of \mathcal{A} of type (1). Now assume that the inequality $(-\alpha_1, \dots, -\alpha_m)^T \nu \leq 0$ is binding for some normalized X -circuit ν of \mathcal{A} . Since the sign pattern $(-, +, +)$ for (ν_i, ν_j, ν_k) in conjunction with $\mathbf{1}^T \nu = 0$ leads to $(-\alpha_1, \dots, -\alpha_m)^T \nu < 0$, and the sign pattern $(+, +, -)$ contradicts the X -circuit condition $\sigma_X(-\mathcal{A}\nu) < \infty$, we obtain the normalized X -circuits of \mathcal{A} of type (2).

For the example classes of the nonnegative orthant and the cube $[-1, 1]^n$, we refer the reader to [27].

4. SUBLINEAR CIRCUITS IN AGE CONES

In this section, we show how the X -AGE cones $C_X(\mathcal{A}, \beta)$ can be further decomposed using sublinear circuits. These decompositions lay the foundation to understand the extreme rays of the conditional SAGE cone $C_X(\mathcal{A})$. Our first result here is a necessary criterion for an X -AGE function f to be extremal in $C_X(\mathcal{A}, \beta)$, which states that all of its relative entropy certificates must be X -circuits (see Theorem 4.2). Definition 4.3 introduces λ -witnessed AGE cones as the subset of signomials in $C_X(\mathcal{A}, \beta)$ whose nonnegativity is certified by a given normalized vector λ . Theorem 4.4 then decomposes $C_X(\mathcal{A}, \beta)$ through the λ -witnessed AGE cones, where λ is a normalized X -circuit. As a consequence, for polyhedral X , the cone $C_X(\mathcal{A}, \beta)$ is power-cone representable (see Corollary 4.5).

In the last part of this section we prove two propositions on explicit representations for primal and dual λ -witnessed AGE cones. Proposition 4.7 in particular is very important for a characterization of dual X -SAGE cones, as it reveals a multiplicative convexity property used extensively in Section 5.

The following lemma provides a construction to decompose an X -AGE function into simpler summands, under a local linearity condition on the support function $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$.

Lemma 4.1. *Let $f = \sum_{\alpha \in \mathcal{A}} c_\alpha e^\alpha$ be X -AGE with negative term $c_\beta < 0$. If ν is a relative entropy certificate for f which can be written as a convex combination $\nu = \sum_{i=1}^k \theta_i \nu^{(i)}$ of non-proportional $\nu^{(i)} \in N_\beta$ and $\tilde{\nu} \mapsto \sigma_X(-\mathcal{A}\tilde{\nu})$ is linear on $\text{conv}\{\nu^{(i)}\}_{i=1}^k$, then f is not extremal in $C_X(\mathcal{A}, \beta)$.*

Proof. Construct vectors $c^{(i)}$ by

$$c_\alpha^{(i)} = \begin{cases} (c_\alpha/\nu_\alpha)\nu_\alpha^{(i)} & \text{if } \alpha \in \nu^+ \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } \alpha \in \mathcal{A} \setminus \beta, \quad (7)$$

and $c_\beta^{(i)} = \sigma_X(-\mathcal{A}\nu^{(i)}) + D(\nu_{\setminus\beta}^{(i)}, ec_{\setminus\beta}^{(i)})$. These $c^{(i)}$ define X -AGE signomials by construction, and they inherit non-proportionality from the $\nu^{(i)}$. We need to show that $\sum_{i=1}^k \theta_i c^{(i)} \leq c$, which will establish that f can be decomposed as a sum of these non-proportional X -AGE functions (possibly with an added posynomial).

For indices $\alpha \in \nu^+$, the construction (7) relative to ν and $\{\nu^{(i)}\}_{i=1}^k$ actually ensures $\sum_{i=1}^k \theta_i c_\alpha^{(i)} = c_\alpha$. For indices $\alpha \in \text{supp } c \setminus \text{supp } \nu$ we have $\sum_{i=1}^k \theta_i c_\alpha^{(i)} = 0 \leq c_\alpha$. The definitions of $\nu^{(i)}$ ensure

$$\sigma_X(-\mathcal{A}\nu) = \sigma_X\left(-\mathcal{A}\left(\sum_{i=1}^k \theta_i \nu^{(i)}\right)\right) = \sum_{i=1}^k \theta_i \sigma_X(-\mathcal{A}\nu^{(i)}). \quad (8)$$

Meanwhile, (7) provides $\nu_\alpha^{(i)}/c_\alpha^{(i)} = \nu_\alpha/c_\alpha$, which may be combined with $\sum_{i=1}^k \theta_i \nu_\alpha^{(i)} = \nu_\alpha \forall \alpha \in \mathcal{A}$ to deduce

$$\sum_{i=1}^k \theta_i D(\nu_{\setminus\beta}^{(i)}, ec_{\setminus\beta}^{(i)}) = D(\nu_{\setminus\beta}, ec_{\setminus\beta}). \quad (9)$$

We combine (8) and (9) to obtain the desired result

$$\sum_{i=1}^k \theta_i c_\beta^{(i)} = \sum_{i=1}^k \theta_i \left(\sigma_X(-\mathcal{A}\nu^{(i)}) + D(\nu_{\setminus\beta}^{(i)}, ec_{\setminus\beta}^{(i)}) \right) = \sigma_X(-\mathcal{A}\nu) + D(\nu_{\setminus\beta}, ec_{\setminus\beta}) \leq c_\beta.$$

□

Theorem 4.2. *Let $f = \sum_{\alpha \in \mathcal{A}} c_\alpha e^\alpha$ be X -AGE with negative term $c_\beta < 0$. If f has a relative entropy certificate which is not an X -circuit, then f is not extremal in $C_X(\mathcal{A}, \beta)$.*

Proof. If f is an X -AGE function with $c_\beta < 0$ and ν satisfies (3), then we must have $\nu \neq \mathbb{0}$ and $\sigma_X(-\mathcal{A}\nu) < \infty$. By the definition of an X -circuit, ν may be written as a convex combination $\nu = \theta\nu^{(1)} + (1-\theta)\nu^{(2)}$ where $\bar{\nu} \mapsto \sigma_X(-\mathcal{A}\bar{\nu})$ is linear on $[\nu^{(1)}, \nu^{(2)}]$, and furthermore the $\nu^{(i)}$ are not proportional. We can therefore invoke Lemma 4.1 to prove the claim. □

In the remainder of this section we eliminate the degree of freedom associated with ν laying on a ray. For each $\beta \in \mathcal{A}$, we introduce the following notation for the associated set of normalized X -circuits of \mathcal{A}

$$\Lambda_X(\mathcal{A}, \beta) = \{\lambda \in N_\beta : \lambda \text{ is an } X\text{-circuit of } \mathcal{A}, \lambda_\beta = -1\}.$$

The set of all normalized X -circuits of \mathcal{A} is denoted $\Lambda_X(\mathcal{A})$. The main reason for introducing this notation is how it interacts with the following definition.

Definition 4.3. Given a vector $\lambda \in N_\beta$ with $\lambda_\beta = -1$, the λ -witnessed AGE cone is

$$C_X(\mathcal{A}, \lambda) = \left\{ \sum_{\alpha \in \mathcal{A}} c_\alpha e^\alpha : \prod_{\alpha \in \lambda^+} \left[\frac{c_\alpha}{\lambda_\alpha} \right]^{\lambda_\alpha} \geq -c_\beta \exp(\sigma_X(-\mathcal{A}\lambda)), c_{\setminus\beta} \geq \mathbb{0} \right\}. \quad (10)$$

We show below that every signomial in $C_X(\mathcal{A}, \lambda)$ is nonnegative on X . The term “witnessed” in “ λ -witnessed AGE cone” is chosen to reflect the defining role of λ in the nonnegativity certificate. We only use λ -witnessed AGE cones for theoretical purposes, and only with $\lambda \in \Lambda_X(\mathcal{A})$. Possible computational uses (particularly with $\lambda \notin \Lambda_X(\mathcal{A})$) are offered in Section 7.

Theorem 4.4. *Let $\Lambda_X(\mathcal{A}) \neq \emptyset$. The cone $C_X(\mathcal{A}, \beta)$ can be written as the convex hull of λ -witnessed AGE cones, where λ runs over the normalized X -circuits, that is,*

$$C_X(\mathcal{A}, \beta) = \text{conv} \bigcup_{\lambda \in \Lambda_X(\mathcal{A}, \beta)} C_X(\mathcal{A}, \lambda).$$

Note here that for any $\beta \in \mathcal{A}$ and (normalized) $\lambda \in N_\beta$, we have $\mathbb{R}_+^{\mathcal{A}} \subset C_X(\mathcal{A}, \lambda)$.

Proof. Theorem 4.2 already tells us that for $\Lambda_X(\mathcal{A}) \neq \emptyset$, $C_X(\mathcal{A}, \beta)$ may be expressed as the convex hull of X -AGE functions $f = \sum_{\alpha \in \mathcal{A}} c_\alpha e^\alpha$ which have X -circuits as relative entropy certificates. Therefore it suffices to show that (i) for any such function, the normalized X -circuit $\lambda = \nu/|\nu_\beta|$ is such that (c, λ) satisfy the condition in (10), and (ii) if *any* (c, λ) satisfy (10), then the resulting signomial is nonnegative on X . We will actually do both of these in one step.

Suppose $\nu \in N_\beta$ is restricted to satisfy $\nu = s\lambda$ for a variable $s \geq 0$ and a fixed $\lambda \in \Lambda_X(\mathcal{A}, \beta)$. It suffices to show that the set of $c \in \mathbb{R}^{\mathcal{A}}$ for which

$$\exists s \geq 0 : \nu = s\lambda \text{ and } \sigma_X(-\mathcal{A}\nu) + D(\nu_{\setminus\beta}, ec_{\setminus\beta}) \leq c_\beta$$

is the same as (10).

Let $r(\nu) = \sigma_X(-\mathcal{A}\nu) + D(\nu_{\setminus\beta}, ec_{\setminus\beta})$. Apply positive homogeneity of the support function to see $\sigma_X(-\mathcal{A}\nu) = |\nu_\beta| \sigma_X(\mathcal{A}\nu/|\nu_\beta|)$, and use $\nu = s\lambda$ to infer $s = |\nu_\beta|$ and $\sigma_X(-\mathcal{A}\nu/|\nu_\beta|) = \sigma_X(-\mathcal{A}\lambda)$. Abbreviate $d := \sigma_X(-\mathcal{A}\lambda)$ and substitute $\sum_{\alpha \in \lambda^+} \nu_\alpha = |\nu_\beta|$ to obtain

$$r(\nu) = \sum_{\alpha \in \lambda^+} (\nu_\alpha \log(\nu_\alpha/c_\alpha) - \nu_\alpha + \nu_\alpha d).$$

The term d may be moved into the logarithm by identifying $\nu_\alpha d = \nu_\alpha \log(1/\exp(-d))$. For $\alpha \in \lambda^+$ we define scaled terms $\tilde{c}_\alpha = c_\alpha \exp(-d)$, so that $r(\nu) = \sum_{\alpha \in \lambda^+} \nu_\alpha \log(\nu_\alpha/\tilde{c}_\alpha) - \nu_\alpha$. By Proposition 8.1, there exists a $\nu = s\lambda$ for which $r(\nu) \leq c_\beta$ if and only if

$$-c_\beta \leq \prod_{\alpha \in \lambda^+} [\tilde{c}_\alpha/\lambda_\alpha]^{\lambda_\alpha}. \quad (11)$$

Since $[\tilde{c}_\alpha/\lambda_\alpha]^{\lambda_\alpha} = [c_\alpha/\lambda_\alpha]^{\lambda_\alpha} (\exp(-d))^{\lambda_\alpha}$ and $\prod_{\alpha \in \lambda^+} (\exp(-d))^{\lambda_\alpha} = \exp(-d)$, (11) can be recognized as the inequality occurring within (10), which completes the proof. \square

Theorem 4.4 shows how λ -witnessed AGE cones provide a window to the structure of *full* AGE cones $C_X(\mathcal{A}, \beta)$. To appreciate the benefit of this perspective, it is necessary to consider the more elementary “power cone.” In our context, the primal power cone associated with a normalized X -circuit $\lambda \in \mathbb{R}^{\mathcal{A}}$ is

$$\text{Pow}(\lambda) = \{z \in \mathbb{R}^{\text{supp } \lambda} : \prod_{\alpha \in \lambda^+} z_\alpha^{\lambda_\alpha} \geq |z_\beta|, z_{\setminus\beta} \geq \mathbf{0}, \beta := \lambda^-\};$$

the corresponding dual cone is given by

$$\text{Pow}(\lambda)^* = \{w \in \mathbb{R}^{\text{supp } \lambda} : \prod_{\alpha \in \lambda^+} [w_\alpha/\lambda_\alpha]^{\lambda_\alpha} \geq |w_\beta|, w_{\setminus\beta} \geq \mathbf{0}, \beta := \lambda^-\}.$$

It should be evident that $C_X(\mathcal{A}, \lambda)$ can be formulated in terms of a dual λ -weighted power cone; a precise formula is provided momentarily. For now we give a corollary concerning power cone representability and second-order representability of $C_X(\mathcal{A})$ when X is a polyhedron (see [2, 3] for formal definitions).

Corollary 4.5. *If X is a polyhedron, then $C_X(\mathcal{A})$ is power cone representable. If in addition $\mathcal{A}^T X$ is rational, then $C_X(\mathcal{A})$ is second-order representable and thus has semidefinite extension degree 2.*

Proof. We can assume $\Lambda_X(\mathcal{A}) \neq \emptyset$, since otherwise $C_X(\mathcal{A}) = \mathbb{R}_+^A$ and the claim follows. By Theorem 3.7, polyhedral X have finitely many X -circuits, up to scaling. Apply Theorem 4.2 and finiteness of the normalized circuits $\Lambda_X(\mathcal{A})$ to write

$$C_X(\mathcal{A}) = \sum_{\lambda \in \Lambda_X(\mathcal{A})} C_X(\mathcal{A}, \lambda).$$

The first claim follows as each of the finitely many sets $C_X(\mathcal{A}, \lambda)$ appearing in the above sum are (dual) power cone representable. For the second claim observe that under the rationality assumptions we have $\Lambda_X(\mathcal{A}) \subset \mathbb{Q}^A$. Using $\beta := \lambda^-$ and $m := |\text{supp } \lambda|$, it is known that the m -dimensional λ -weighted power cone (and its dual) are second-order representable when $\lambda_{\setminus \beta}$ is a rational vector in the $(m - 1)$ -dimensional probability simplex [3, Section 3.4]. The last claim follows as the semidefinite extension degree of the second-order cone is two [3, Section 2.3]. \square

The first part of Corollary 4.5 generalizes the case $X = \mathbb{R}^n$ considered by Papp for polynomials [33]. That aspect of the corollary has uses in computational optimization when applied judiciously. The second part of Corollary 4.5 generalizes results by Averkov [2] and Wang and Magron [41] for ordinary SAGE polynomials, and recent results by Naumann and Theobald for several types of ordinary SAGE-like certificates [26]. We have deliberately framed the second part of the corollary in abstract terms (semidefinite extension degree), because that aspect of the corollary seems not useful for computational optimization.

We now work towards finding a simple representation of dual λ -witnessed AGE cones $C_X(\mathcal{A}, \lambda)^*$. We begin this process by regarding the primal as a cone of coefficients contained in \mathbb{R}^A , and finding an explicit representation of the primal in terms of the elementary dual power cone $\text{Pow}(\lambda)^*$. Towards that end we introduce a diagonal linear operator $S_\lambda : \mathbb{R}^A \rightarrow \mathbb{R}^{\text{supp } \lambda}$ where $(S_\lambda w)_\alpha = w_\alpha$ for $\alpha \in \lambda^+$, and $(S_\lambda w)_\beta = w_\beta \exp(\sigma_X(-\mathcal{A}\lambda))$ for $\beta := \lambda^-$. Recall that $\delta_\beta \in \mathbb{R}^A$ denotes the standard basis vector corresponding to $\beta \in \mathcal{A}$, i.e., $\delta_\beta^T w = w_\beta$ for $w \in \mathbb{R}^A$.

Proposition 4.6. *For $\lambda \in N_\beta$ with $\lambda_\beta = -1$ and $\sigma_X(-\mathcal{A}\lambda) < \infty$, the λ -witnessed AGE cone admits the representation*

$$C_X(\mathcal{A}, \lambda) = \{c \in \mathbb{R}^A : \beta := \lambda^-, c_{\setminus \beta} \geq \mathbb{0}, (S_\lambda c - r\delta_\beta) \in \text{Pow}(\lambda)^*, r \geq 0\}. \quad (12)$$

Proof. First, we note that some inequality constraints $c_{\setminus \beta} \geq \mathbb{0}$ are implied by $(S_\lambda c - r\delta_\beta) \in \text{Pow}(\lambda)^*$. It is necessary to include the inequality constraints explicitly, to account for the case when $\text{supp } \lambda \subsetneq \mathcal{A}$. The condition $(S_\lambda c - r\delta_\beta) \in \text{Pow}(\lambda)^*$ can be rewritten as

$$\prod_{\alpha \in \lambda^+} [c_\alpha / \lambda_\alpha]^{\lambda_\alpha} \geq |c_\beta \exp(\sigma_X(-\mathcal{A}\lambda)) - r|. \quad (13)$$

Meanwhile, the minimum of $|c_\beta \exp(\sigma_X(-\mathcal{A}\lambda)) - r|$ over $r \geq 0$ is attained at $r = 0$ when $c_\beta < 0$ and $r = c_\beta \exp(\sigma_X(-\mathcal{A}\lambda))$ when $c_\beta \geq 0$. In the $c_\beta < 0$ case the constraint (13) becomes

$$\prod_{\alpha \in \lambda^+} [c_\alpha / \lambda_\alpha]^{\lambda_\alpha} \geq -c_\beta \exp(\sigma_X(-\mathcal{A}\lambda)).$$

In the $c_\beta \geq 0$ case the constraint (13) is vacuous, since $\prod_{\alpha \in \lambda^+} [c_\alpha / \lambda_\alpha]^{\lambda_\alpha} \geq 0$ is implied by $c_{\setminus \beta} \geq \mathbf{0}$. As the constraint in the preceding display is similarly vacuous when $c_\beta > 0$, we see that it can be used in lieu of (13) without loss of generality. \square

We can appeal to Proposition 4.6 to find a representation for $C_X(\mathcal{A}, \lambda)^*$ which is analogous to Equation (10). Again, the dual is computed by regarding the primal as a cone of *coefficients*.

Proposition 4.7. *For $\lambda \in N_\beta$ with $\lambda_\beta = -1$ and $\sigma_X(-\mathcal{A}\lambda) < \infty$, the dual λ -witnessed AGE cone is given by*

$$C_X(\mathcal{A}, \lambda)^* = \left\{ v \in \mathbb{R}_+^A : \beta := \lambda^-, \exp(\sigma_X(-\mathcal{A}\lambda)) \prod_{\alpha \in \lambda^+} v_\alpha^{\lambda_\alpha} \geq v_\beta \right\}. \quad (14)$$

Proof. Let $\beta = \lambda^-$ as is usual. To $v \in \mathbb{R}^A$ associate $\text{Val}(v) = \inf\{v^T c : c \in C_X(\mathcal{A}, \lambda)\}$. A vector v belongs to $C_X(\mathcal{A}, \lambda)^*$ if and only if $\text{Val}(v) = 0$. We will find constraints on v so that the dual feasible set for computing $\text{Val}(v)$ is nonempty, which in turn will imply $\text{Val}(v) = 0$.

We begin by noting that for any element $\alpha \in \mathcal{A} \setminus \text{supp } \lambda$, the only constraints on c_α, v_α for $c \in C_X(\mathcal{A}, \lambda), v \in C_X(\mathcal{A}, \lambda)^*$ are $c_\alpha \geq 0, v_\alpha \geq 0$; therefore we assume $\mathcal{A} = \text{supp } \lambda$ for the remainder of the proof. When considering the given expression for $\text{Val}(v)$ as a primal problem, we compute a dual using (12) from Proposition 4.6. Under the assumption $\mathcal{A} = \text{supp } \lambda$, the constraint $c_{\setminus \beta} \geq \mathbf{0}$ is implied by $(S_\lambda c - r\delta_\beta) \in \text{Pow}(\lambda)^*$. Therefore when forming a Lagrangian for $\text{Val}(v)$ using (12), the dual variable to “ $c_{\setminus \beta} \geq \mathbf{0}$ ” may be omitted.

For the remaining constraints $(S_\lambda c - r\delta_\beta) \in \text{Pow}(\lambda)^*$ and $r \geq 0$ we use dual variables $\mu \in \text{Pow}(\lambda)$ and $t \in \mathbb{R}_+$ respectively; the Lagrangian is

$$\begin{aligned} \mathcal{L}(c, r, \mu, t) &= v^T c - \mu^T (S_\lambda c - r\delta_\beta) - tr \\ &= c^T (v - S_\lambda^T \mu) - r(t - \mu_\beta). \end{aligned}$$

For the Lagrangian to be bounded below over $c \in \mathbb{R}^A$ and $r \in \mathbb{R}$, it is necessary and sufficient that $v = S_\lambda^T \mu$ and $\mu_\beta = t$. Since we have assumed $\text{supp } \lambda = \mathcal{A}$ and $\sigma_X(-\mathcal{A}\lambda) < \infty$, the diagonal linear operator S_λ is symmetric positive definite, so we can express the requirements on μ, t as

$$S_\lambda^{-1} v = \mu \quad \text{and} \quad \mu_\beta = t.$$

Therefore the conditions $S_\lambda^{-1} v \in \text{Pow}(\lambda), v_\beta \geq 0$ are equivalent to

$$\begin{aligned} \text{Val}(v) &= \inf \left\{ \sup \{ \mathcal{L}(c, r, \mu, t) : (\mu, t) \in \text{Pow}(\lambda) \times \mathbb{R}_+ \} : (c, r) \in \mathbb{R}^A \times \mathbb{R} \right\} \\ &= \sup \left\{ \inf \{ \mathcal{L}(c, r, \mu, t) : (c, r) \in \mathbb{R}^A \times \mathbb{R} \} : (\mu, t) \in \text{Pow}(\lambda) \times \mathbb{R}_+ \right\} = 0. \end{aligned}$$

The proposition follows by applying the definitions of $\text{Pow}(\lambda)$ and S_λ . \square

5. REDUCED SUBLINEAR CIRCUITS IN SAGE CONES

The previous section showed that an X -SAGE cone is generated by X -circuits. Here we seek a much sharper characterization: are all X -circuits really necessary? The answer to this question depends on whether one means to reconstruct an individual AGE cone, or the larger SAGE cone. For example, by reinterpreting results from [23], we may infer that every simplicial \mathbb{R}^n -circuit $\lambda \in \Lambda_{\mathbb{R}^n}(\mathcal{A}, \beta)$ generates a λ -witnessed AGE cone containing an extreme ray of $C_{\mathbb{R}^n}(\mathcal{A}, \beta)$. In this way, every \mathbb{R}^n -circuit is needed if one requires complete reconstruction of individual AGE cones. However, Katthän, Naumann, and Theobald showed that many extreme rays of AGE cones are *not* extreme when considered in the sum $C_{\mathbb{R}^n}(\mathcal{A}) = \sum_{\beta \in \mathcal{A}} C_{\mathbb{R}^n}(\mathcal{A}, \beta)$. Specifically, an \mathbb{R}^n -circuit $\lambda \in \Lambda_{\mathbb{R}^n}(\mathcal{A})$ is only needed in $C_{\mathbb{R}^n}(\mathcal{A})$ if exactly one element of \mathcal{A} hits the relative interior of $\text{conv}(\text{supp } \lambda)$ [17, Proposition 4.4]. Circuits satisfying this property were called *reduced*. The goal of this section is to develop a reducedness criterion for X -circuits that yields the most efficient construction of $C_X(\mathcal{A})$ by λ -witnessed AGE cones, see Theorems 5.5 and 5.6. Achieving this goal is more difficult than obtaining the results from earlier sections. Therefore we begin by summarizing and discussing the results, and we provide proofs in later subsections.

5.1. Definitions, results, and discussion. The definition of a reduced \mathbb{R}^n -circuit is of a purely combinatorial nature, involving the circuit's support. This is appropriate because when speaking of affine-linear simplicial circuits, the normalized vector representation λ is completely determined by its support. In the context of X -circuits, we no longer have this property. Therefore when developing reduced X -circuits it is useful to have a different characterization of reduced \mathbb{R}^n -circuits. Here we can consider how Forsgård and de Wolff defined the *Reznick cone of \mathcal{A}* as the conic hull $R_{\mathbb{R}^n}(\mathcal{A}) := \text{co } \Lambda_{\mathbb{R}^n}(\mathcal{A})$ and – in the language of Katthän et al. – subsequently proved that an \mathbb{R}^n -circuit λ is an edge generator of $R_{\mathbb{R}^n}(\mathcal{A})$ if and only if it is reduced [12].

Our definition of reduced X -circuits involves edge generators of a certain cone in one higher dimension than the Reznick cone. To describe the cone and facilitate later analysis, we need the following definition.

Definition 5.1. The *functional form* of an X -circuit $\nu \in \mathbb{R}^{\mathcal{A}}$ is $\phi_\nu : \mathbb{R}^{\mathcal{A}} \rightarrow \mathbb{R}$ defined by

$$\phi_\nu(y) = \sum_{\alpha \in \mathcal{A}} y_\alpha \nu_\alpha + \sigma_X(-\mathcal{A}\nu).$$

We routinely overload notation and use $\phi_\nu = (\nu, \sigma_X(-\mathcal{A}\nu)) \in \mathbb{R}^{\mathcal{A}} \times \mathbb{R}$ to denote the functional form of a given X -circuit. When representing the functional form of an X -circuit by a vector in $\mathbb{R}^{\mathcal{A}} \times \mathbb{R}$, the scalar $\phi_\nu(y)$ can be expressed as an inner product $\phi_\nu(y) = (y, 1)^T \phi_\nu$.

Definition 5.2. The *circuit-generated cone* (shortly, *CG cone*) of (\mathcal{A}, X) is

$$G_X(\mathcal{A}) = \text{co}(\{\phi_\lambda : \lambda \in \Lambda_X(\mathcal{A})\} \cup \{(\mathbf{0}, 1)\}),$$

where $(\mathbf{0}, 1) \in \mathbb{R}^{\mathcal{A}} \times \mathbb{R}$.

The idea of generating a cone from augmented circuit vectors $(\nu, \sigma_X(-\mathcal{A}\nu)) \in \mathbb{R}^{\mathcal{A}} \times \mathbb{R}$ clearly parallels Theorem 3.6. While the cones from Theorem 3.6 are considered for one $\beta \in \mathcal{A}$ at a

time, the CG cone accounts for all X -circuits at once. The CG cone also includes an extra generator that ultimately serves to make the following definition more stringent.

Definition 5.3. The *reduced X -circuits* of \mathcal{A} are the vectors ν where $\nu/\|\nu\|_\infty \in \Lambda_X(\mathcal{A})$ and the corresponding functional form ϕ_ν generates an extreme ray of $G_X(\mathcal{A})$. The set of *normalized reduced X -circuits* is henceforth denoted $\Lambda_X^*(\mathcal{A})$.

There is a subtle issue here that in order for reduced X -circuits to be of any use to us, the CG cone must be pointed (else $G_X(\mathcal{A})$ would have no extreme rays whatsoever). We show later in this section that our stated assumption of linear independence of $\{e^\alpha\}_{\alpha \in \mathcal{A}}$ on X ensures $G_X(\mathcal{A})$ is pointed. Regardless of whether or not the CG cone is pointed, we have the following theorem.

Theorem 5.4. $C_X(\mathcal{A})^* = \text{cl}\{\exp y : (y, 1) \in G_X(\mathcal{A})^*\}$.

Theorem 5.4 is noteworthy in several respects. It demonstrates that $C_X(\mathcal{A})^*$ is convex in the usual sense and convex under a logarithmic transformation $S \mapsto \log S = \{y : \exp y \in S\}$. This second form of convexity is a significant structural property. For example, if we know that the log of the moment cone $\text{cl}(\text{co}\{\exp(\mathcal{A}^T x) : x \in X\})$ is not convex, then it should be that $C_X(\mathcal{A})$ does not contain all X -nonnegative signomials on \mathcal{A} . Additionally, Theorem 5.4 can be reverse-engineered to arrive at the concept of a reduced X -circuit: the definition is chosen so that $(y, 1)$ belongs to $G_X(\mathcal{A})^*$ if and only if $\phi_\lambda(y) \geq 0$ for all λ in $\Lambda_X^*(\mathcal{A})$. Here, Theorem 5.4 is a tool that we combine with convex duality to obtain the following results.

Theorem 5.5. If $\Lambda_X(\mathcal{A})$ is empty, then $C_X(\mathcal{A}) = \mathbb{R}_+^{\mathcal{A}}$. Otherwise,

$$C_X(\mathcal{A}) = \text{cl}\left(\text{conv}\bigcup\left\{C_X(\mathcal{A}, \lambda) : \lambda \in \Lambda_X^*(\mathcal{A})\right\}\right). \quad (15)$$

We point out how Theorem 5.5 involves a closure around the union over λ -witnessed AGE cones, while Theorem 4.4 has no such closure. The need for the closure here stems from an application of an infinite version of conic duality in the course of the theorem's proof, while our proof of Theorem 4.4 required no duality at all. The requisite use of conic duality is simpler when X is a polyhedron, as the following theorem suggests.

Theorem 5.6. If X is a polyhedron and $\Lambda_X(\mathcal{A})$ is nonempty, then the associated conditional SAGE cone is given by the finite Minkowski sum

$$C_X(\mathcal{A}) = \sum_{\lambda \in \Lambda_X^*(\mathcal{A})} C_X(\mathcal{A}, \lambda). \quad (16)$$

Moreover, there is no proper subset $\Lambda \subsetneq \Lambda_X^*(\mathcal{A})$ for which $C_X(\mathcal{A}) = \sum_{\lambda \in \Lambda} C_X(\mathcal{A}, \lambda)$.

The first part of Theorem 5.6 follows easily from the arguments we use to prove Theorem 5.5. The second part of the theorem is much more delicate, and in fact is the reason why $G_X(\mathcal{A})$ is defined in the manner of 5.2, rather than merely $\text{co}\{\phi_\lambda : \lambda \in \Lambda_X(\mathcal{A})\}$.

The task of actually finding the reduced X -circuits of \mathcal{A} is difficult. When X is a polyhedron there are finitely many such X -circuits, but the naive method for finding them involves Fourier-Motzkin elimination on a set of potentially very high dimension. There is more hope for this problem when X is a cone. In that case, X -circuits are the extreme rays of $(\ker \mathcal{A} + \mathcal{A}^\dagger X^*) \cap N_\beta$

for $\beta \in \mathcal{A}$, and no lifting is needed to find these extreme rays with a computer. The reduced X -circuits could then be computed by finding the extreme rays of the convex cone generated by the X -circuits. The following detailed example finds the reduced X -circuits of \mathcal{A} in the univariate case with $X = [0, \infty)$. The claim made in the example is used in Section 6.

Example 5.7. We continue the running example of $X = [0, \infty)$ from Example 3.8. In particular recall $\mathcal{A} = \{\alpha_1, \dots, \alpha_m\}$ for $\alpha_1 < \dots < \alpha_m$, indexing by $i \in [m]$, and working with standard basis $\delta_i \in \mathbb{R}^m$. We claim that

$$\Lambda_{[0, \infty)}^*(\mathcal{A}) = \{\delta_2 - \delta_1\} \cup \Lambda_{\mathbb{R}}^*(\mathcal{A}) \quad (17)$$

where we have the following formula from [17, Prop. 4.4]

$$\Lambda_{\mathbb{R}}^*(\mathcal{A}) = \left\{ \left(\frac{\alpha_{i+1} - \alpha_i}{\alpha_{i+1} - \alpha_{i-1}} \right) \delta_{i-1} + \left(\frac{\alpha_i - \alpha_{i-1}}{\alpha_{i+1} - \alpha_{i-1}} \right) \delta_{i+1} - \delta_i : 1 < i < m \right\}.$$

As a first step towards seeing this, observe that since $X = [0, \infty)$ is a cone, the functional form of a $[0, \infty)$ -circuit ν is simply $\phi_\nu(y) = \sum_{i=1}^m y_i \nu_i$. Hence, the reduced $[0, \infty)$ -circuits are exactly the edge generators of the cone $\text{co } \Lambda_{[0, \infty)}$ generated by all the $[0, \infty)$ -circuits of types (1) and (2) listed in Example 3.8. Therefore, we have to show that $\{\delta_2 - \delta_1\} \cup \Lambda_{\mathbb{R}}^*(\mathcal{A})$ are exactly the normalized edge generators of $\text{co } \Lambda_{[0, \infty)}$.

For the X -circuits $\delta_j - \delta_i$ ($j > i$) of type (1) in Example 3.8, we show they decompose if $j > i + 1$ or $i > 1$. For $j > i + 1$, this is apparent from the decomposition

$$\delta_j - \delta_i = (\delta_j - \delta_{j-1}) + (\delta_{j-1} - \delta_i).$$

For $j = i + 1$ and $i > 1$, we can use the decomposition

$$\delta_{i+1} - \delta_i = \left(-\frac{\alpha_{i+1} - \alpha_i}{\alpha_i - \alpha_{i-1}} \delta_{i-1} + \frac{\alpha_{i+1} - \alpha_i}{\alpha_i - \alpha_{i-1}} \delta_i \right) + \left(\frac{\alpha_{i+1} - \alpha_i}{\alpha_i - \alpha_{i-1}} \delta_{i-1} - \frac{\alpha_{i+1} - \alpha_{i-1}}{\alpha_i - \alpha_{i-1}} \delta_i + \delta_{i+1} \right)$$

into X -circuits with three non-vanishing components. As final consideration for type (1), the X -circuit $\delta_2 - \delta_1$ cannot be written as a conic combination of X -circuits with three non-zero entries, because any conic combination of those X -circuits has a positive entry in its non-vanishing component with maximal index. For X -circuits of type (2) from Example 3.8, simply note that these are also \mathbb{R} -circuits. Therefore a necessary condition for a type (2) X -circuit λ to be extremal in $\text{co } \Lambda_{[0, \infty)}$ is that λ belongs to $\Lambda_{\mathbb{R}}^*(\mathcal{A})$.

It remains to show that none of the remaining X -circuits can be written as a convex combination of the others. First note that an X -circuit $\nu \in \Lambda_{\mathbb{R}}^*(\mathcal{A})$ cannot be decomposed into a sum which involves an X -circuit $\tilde{\nu}$ with two non-vanishing components. Namely, since $\mathcal{A}\nu = 0$ and $\mathcal{A}\tilde{\nu} > 0$, we would obtain for the other summand $\nu - \tilde{\nu}$ the property $\mathcal{A}(\nu - \tilde{\nu}) < 0$ and thus $\sigma_{[0, \infty)}(-\mathcal{A}(\nu - \tilde{\nu})) = \infty$, a contradiction. And of course it is trivially true that no element $\lambda \in \Lambda_{\mathbb{R}}^*(\mathcal{A})$ can be written as a convex combination of other such elements. Since $\text{co } \Lambda_{[0, \infty)}$ is finitely generated and there is no $S \subsetneq \{\delta_2 - \delta_1\} \cup \Lambda_{\mathbb{R}}^*(\mathcal{A})$ for which $\text{co } \Lambda_{[0, \infty)} = \text{co } S$, we conclude that $\{\delta_2 - \delta_1\} \cup \Lambda_{\mathbb{R}}^*(\mathcal{A})$ are the reduced X -circuits of \mathcal{A} .

The remainder of this section is organized as follows. Section 5.2 proves Theorem 5.4, which is instrumental in later subsections. In Section 5.3 we introduce and prove a certain representation

result for the CG cone. Given the groundwork laid in these two subsections, Section 5.4 proves Theorem 5.5 in very short order. Section 5.5 proves Theorem 5.6 by refining the arguments from Section 5.4.

5.2. **Proof of Theorem 5.4.** We begin with the following simple lemma.

Lemma 5.8. *If $S \subset T$ are convex sets where S is closed and $S \cap \text{ri } T \neq \emptyset$, then $S = \text{cl}(S \cap \text{ri } T)$.*

Proof. Rockafellar's [37, Theorem 18.2] states that every relatively open set contained in T is contained in the relative interior of some face of T . By our assumption $S \cap \text{ri } T \neq \emptyset$, the only face of T which contains S is T itself. Since $\text{ri } S$ is obviously relatively open, we have $\text{ri } S \subset \text{ri } T$, and the claim follows by the identity $S = \text{cl } \text{ri } S$ for closed convex sets. \square

Proof of Theorem 5.4. Use Rockafellar's [37, Corollary 16.5.2] to invoke Theorem 4.4 from a dual point of view, which gives $C_X(\mathcal{A}, \beta)^* = \bigcap C_X(\mathcal{A}, \lambda)^*$, where the intersection runs over all $\lambda \in \Lambda_X(\mathcal{A}, \beta)$. Then Proposition 4.7 implies

$$C_X(\mathcal{A})^* = \left\{ v \in \mathbb{R}_+^A : \forall \lambda \in \Lambda_X(\mathcal{A}), \beta := \lambda^-, \exp(\sigma_X(-\mathcal{A}\lambda)) \prod_{\alpha \in \lambda^+} v_\alpha^{\lambda_\alpha} \geq v_\beta \right\}. \quad (18)$$

We claim that $C_X(\mathcal{A})^*$ can be represented as the closure of its intersection with the positive orthant, that is, $C_X(\mathcal{A})^* = \text{cl}(C_X(\mathcal{A})^* \cap \mathbb{R}_{++}^A)$. Since $C_X(\mathcal{A})$ contains all posynomials and is contained in the nonnegativity cone, the dual $C_X(\mathcal{A})^*$ contains the moment cone but is still contained in the nonnegative orthant. As we have assumed X is nonempty, $C_X(\mathcal{A})^*$ must contain a point $\exp(\mathcal{A}^T x) \in \mathbb{R}_{++}^A$, so $C_X(\mathcal{A})^* \cap \text{ri } \mathbb{R}_+^A \neq \emptyset$. Applying Lemma 5.8 with $S = C_X(\mathcal{A})^*$ and $T = \mathbb{R}_+^A$ gives $C_X(\mathcal{A})^* = \text{cl}(C_X(\mathcal{A})^* \cap \text{ri } \mathbb{R}_+^A) = \text{cl}(C_X(\mathcal{A})^* \cap \mathbb{R}_{++}^A)$.

When considering $C_X(\mathcal{A})^*$ only over the positive orthant, the inequalities

$$\exp(\sigma_X(-\mathcal{A}\lambda)) \prod_{\alpha \in \lambda^+} v_\alpha^{\lambda_\alpha} \geq v_\beta$$

appearing in (18) may be rewritten as

$$\sum_{\alpha \in \lambda^+} \lambda_\alpha \log v_\alpha - \log v_\beta + \sigma_X(-\mathcal{A}\lambda) \equiv \phi_\lambda(y) \geq 0,$$

where we used $\lambda_\beta = -1$ and $y = \log v \in \mathbb{R}^A$. Hence,

$$\begin{aligned} C_X(\mathcal{A})^* &= \text{cl}\{\exp(y) : \phi_\lambda(y) \geq 0 \ \forall \lambda \in \Lambda_X(\mathcal{A})\} \\ &= \text{cl}\{\exp(y) : (y, 1)^T(\lambda, \tau) \geq 0 \ \forall \lambda \in \Lambda_X(\mathcal{A}), \tau \geq \sigma_X(-\mathcal{A}\lambda)\} \\ &= \text{cl}\{\exp(y) : (y, 1)^T(\nu, \tau) \geq 0 \ \forall (\nu, \tau) \in G_X(\mathcal{A})\}. \end{aligned}$$

By the definition of the dual cone from convex analysis, the property $(y, 1)^T(\nu, \tau) \geq 0 \ \forall (\nu, \tau) \in G_X(\mathcal{A})$ is the same as $(y, 1) \in G_X(\mathcal{A})^*$. This completes the proof. \square

The ability to represent $C_X(\mathcal{A})^*$ in terms of $G_X(\mathcal{A})^*$ is key to our proofs of Theorems 5.5 and 5.6. Note that the theorem remains true when $G_X(\mathcal{A})$ is replaced by the smaller set $\text{co}\{\phi_\lambda : \lambda \in \Lambda_X(\mathcal{A})\}$, because the term $(0, 1)$ simply requires $(y, t) \in G_X(\mathcal{A})^*$ to have $t \geq 0$.

5.3. Topological properties of the CG cone. We need some topological properties of the CG cone from Definition 5.2.

Theorem 5.9. $G_X(\mathcal{A}) = \text{co}(\{\phi_\lambda : \lambda \in \Lambda_X^*(\mathcal{A})\} \cup \{(\mathbf{0}, 1)\})$.

The proof of this theorem essentially reduces to showing that $G_X(\mathcal{A})$ is pointed and closed. The pointedness of the CG cone is easy to show, but closedness is a more delicate matter. In fact – our proof that $G_X(\mathcal{A})$ is closed relies on the fact that it is pointed. We therefore prove pointedness before discussing closedness any further.

Lemma 5.10. *The closure of the CG cone contains no lines.*

Proof. We focus on proving $G_X(\mathcal{A})^*$ is full-dimensional. Let $|\mathcal{A}| = m$. We assumed at the outset of the article that the moment cone $M_X(\mathcal{A}) := \text{co}\{\exp(\mathcal{A}^T x) : x \in X\}$ was full-dimensional, i.e., $\dim M_X(\mathcal{A}) = m$; we use that assumption in this lemma. Specifically, since $C_X(\mathcal{A})$ is contained within the nonnegativity cone, we have that $M_X(\mathcal{A}) \subset C_X(\mathcal{A})^*$ and so $\dim C_X(\mathcal{A})^* = m$. By Theorem 5.4 and continuity of the exponential function, we see that if $\dim C_X(\mathcal{A})^* = m$, then the preimage $S := \{y : (y, 1) \in G_X(\mathcal{A})^*\}$ likewise has dimension m . Consider the induced cone associated with S :

$$\text{indco } S = \text{cl}\{(y, t) : t > 0, y/t \in S\} = \text{cl}\{(y, t) : t > 0, (y, t) \in G_X(\mathcal{A})^*\}.$$

The rightmost expression in the above display tells us $\text{indco } S \subset G_X(\mathcal{A})^*$. We claim without proof that since S is a full-dimensional convex set, $\text{indco } S$ is similarly full-dimensional. Taking this claim as given, $\text{indco } S \subset G_X(\mathcal{A})^*$ implies $G_X(\mathcal{A})^*$ is full-dimensional. Because $G_X(\mathcal{A})^*$ is full-dimensional, $\text{cl } G_X(\mathcal{A}) = G_X(\mathcal{A})^{**} \supset G_X(\mathcal{A})$ contains no lines. \square

In the special case where X is a polyhedron, closedness of $G_X(\mathcal{A})$ follows from Theorem 3.7, which tells us that $\Lambda_X(\mathcal{A})$ is finite. To prove closedness for arbitrary convex sets X we need to more carefully appeal to properties of the generating set $\{\phi_\lambda : \lambda \in \Lambda_X(\mathcal{A})\} \cup \{(\mathbf{0}, 1)\}$.

Lemma 5.11. *The CG cone is closed.*

Proof. Let $S_\beta = \{(\lambda, \sigma_X(-\mathcal{A}\lambda)) : \lambda \in \Lambda_X(\mathcal{A}, \beta)\}$. By Theorem 3.6, the elements $\phi_\lambda \in S_\beta$ are edge generators for the closed convex cone $T_\beta = \text{co}\{(\nu, \sigma_X(-\mathcal{A}\nu)) : \nu \in N_\beta, \sigma_X(-\mathcal{A}\nu) < \infty\}$. From S_β we form $S'_\beta := \text{conv } S_\beta$, and find S'_β is isomorphic to $S'_\beta = \{\phi_\lambda \in T_\beta : \lambda_\beta = -1\}$. Because S_β is bounded, S'_β is likewise bounded. Because S'_β is a slice of a closed convex cone T_β , we have that S'_β is closed. Therefore we conclude S'_β is compact.

Now define $S' = (\bigcup_{\beta \in \mathcal{A}} S'_\beta) \cup \{(\mathbf{0}, 1)\}$. The set S' is a compact generating set for $G_X(\mathcal{A})$ which does not contain the origin. Since $\text{cl } G_X(\mathcal{A})$ is known to contain no lines (Lemma 5.10), we apply Proposition 8.2 to S' , $\text{co } S'$ to infer that $\text{co } S' = G_X(\mathcal{A})$ is closed. \square

Proof of Theorem 5.9. Lemmas 5.10 and 5.11 show $G_X(\mathcal{A})$ is closed and pointed. By [37, Corollary 18.5.2], we have that $G_X(\mathcal{A})$ may be expressed as the conic hull of any set of vectors containing all of its extreme rays. Since $S = \{\phi_\lambda : \lambda \in \Lambda_X(\mathcal{A})\} \cup \{(\mathbf{0}, 1)\}$ is a generating set for $G_X(\mathcal{A})$, it must contain all extreme rays of $G_X(\mathcal{A})$. However, by definition of $\Lambda_X^*(\mathcal{A})$, if λ does not belong to $\Lambda_X^*(\mathcal{A})$, then $\phi_\lambda \in S$ does not generate an extreme ray of $G_X(\mathcal{A})$. We may therefore form $T = S \setminus \{\phi_\lambda : \lambda \notin \Lambda_X^*(\mathcal{A})\}$ and still find $G_X(\mathcal{A}) = \text{co } T$. This proves the theorem. \square

5.4. Proof of Theorem 5.5.

Proof of Theorem 5.5. Using the representation $G_X(\mathcal{A}) = \text{co}(\{\phi_\lambda : \lambda \in \Lambda_X^*(\mathcal{A})\} \cup \{(\mathbf{0}, 1)\})$ provided by Theorem 5.9, we can express

$$(y, 1) \in G_X(\mathcal{A})^* \Leftrightarrow (y, 1)^T(\lambda, \sigma_X(-\mathcal{A}\lambda)) \geq 0 \quad \forall \lambda \in \Lambda_X^*(\mathcal{A}). \quad (19)$$

We obtain the following refinement of Equation (18), by combining (19) with Theorem 5.4:

$$C_X(\mathcal{A})^* = \left\{ v \in \mathbb{R}_+^{\mathcal{A}} : \forall \lambda \in \Lambda_X^*(\mathcal{A}), \beta := \lambda^-, \exp(\sigma_X(-\mathcal{A}\lambda)) \prod_{\alpha \in \lambda^+} v_\alpha^{\lambda_\alpha} \geq v_\beta \right\}. \quad (20)$$

Of course, Equation (20) can be written as $C_X(\mathcal{A})^* = \bigcap_{\lambda \in \Lambda_X^*(\mathcal{A})} C_X(\mathcal{A}, \lambda)^*$. We appeal to conic duality principles (again, [37, Corollary 16.5.2]) to obtain the claim of the theorem. \square

5.5. Proof of Theorem 5.6. A conceptual message from the last section is that it can be very useful to analyze $C_X(\mathcal{A})$ in terms of the vectors y where $\exp y$ belongs to $C_X(\mathcal{A})^*$. This section will hammer that message home. We begin with the lemma that ultimately led us to define $G_X(\mathcal{A})$ as per Definition 5.2, rather than as the simpler set $\text{co}\{\phi_\lambda : \lambda \in \Lambda_X(\mathcal{A})\}$.

Lemma 5.12. *If X is polyhedral and $\Lambda \subsetneq \Lambda_X^*(\mathcal{A})$, then there must exist a $\tilde{y} \in \mathbb{R}^{\mathcal{A}}$ satisfying $\phi_{\lambda'}(\tilde{y}) \geq 0$ for all $\lambda' \in \Lambda$, yet for some $\lambda \in \Lambda_X^*(\mathcal{A}) \setminus \Lambda$ we have $\phi_\lambda(\tilde{y}) < 0$.*

Proof. Let $T_1 = \{\phi_\lambda : \lambda \in \Lambda_X^*(\mathcal{A})\} \cup \{(\mathbf{0}, 1)\}$ and $T_2 = \{\phi_\lambda : \lambda \in \Lambda\} \cup \{(\mathbf{0}, 1)\}$. Of course, a vector \tilde{y} satisfies $\phi_{\lambda'}(\tilde{y}) \geq 0$ for all $\lambda' \in \Lambda$ if and only if $(\tilde{y}, 1) \in (\text{co } T_2)^*$. We will show that given the polyhedrality of X and the assumption on Λ , there exists a vector \tilde{y} for which $(\tilde{y}, 1) \in (\text{co } T_2)^* \setminus (\text{co } T_1)^*$. The result will follow since membership of vectors $(y, 1) \in (\text{co } T_1)^*$ is equivalent to $\phi_\lambda(y) \geq 0$ for all $\lambda \in \Lambda_X^*(\mathcal{A})$.

Since X is polyhedral, the cones T_1 and T_2 are also polyhedral (both are finitely generated by Theorem 3.7). Meanwhile, Theorem 5.9 tells us that $G_X(\mathcal{A}) = \text{co } T_1$, and the definition of reduced circuits is such that every $\phi_\lambda \in T_1 \setminus \{(\mathbf{0}, 1)\}$ generates an extreme ray in $G_X(\mathcal{A})$. Since $\Lambda \subsetneq \Lambda_X^*(\mathcal{A})$, there exists a $\phi_\lambda \in T_1 \setminus T_2$ which generates an extreme ray of $G_X(\mathcal{A})$. Therefore $\text{co } T_2$ is a strict subset of $\text{co } T_1 \equiv G_X(\mathcal{A})$. We may take dual cones to find $(\text{co } T_2)^* \supsetneq (\text{co } T_1)^*$. Note that since T_1 and T_2 contain $\{(\mathbf{0}, 1)\}$, the dual cones must be contained in $K = \mathbb{R}^{\mathcal{A}} \times \mathbb{R}_+$. Furthermore, since X is presumed nonempty, Theorem 5.4 tells us there exists a point $(y, 1) \in (\text{co } T_1)^*$, so the relative interiors of $(\text{co } T_1)^*$ and $(\text{co } T_2)^*$ are contained within the relative interior of K . As our last step, use the fact that if one closed polyhedral cone strictly contains another closed polyhedral cone, then there exists a point in the relative interior of the larger cone which may be separated from the smaller cone; apply this to $(\text{co } T_2)^* \supsetneq (\text{co } T_1)^*$ to find a point $(y', t') \in \text{ri}((\text{co } T_2)^*) \setminus (\text{co } T_1)^*$ with $t' > 0$. From this (y', t') we rescale $\tilde{y} = y'/t'$ so that $(\tilde{y}, 1) \in (\text{co } T_2)^* \setminus (\text{co } T_1)^*$. \square

Remark 5.13. We take a moment to unpack the technical dependencies in Lemma 5.12. We explicitly cited Theorem 5.9. Our proof of that result relied on Lemma 5.11, which states that the CG cone is closed, and which we proved by appeal to Theorem 3.6. However, when X is a polyhedron, Lemma 5.11 can alternatively be proven by appeal to Theorem 3.7.

Our next lemma shows how to take a condition stated in terms of Lemma 5.12, and deduce a statement about $C_X(\mathcal{A})^*$. The lemma's proof requires only that X be nonempty and convex.

Lemma 5.14. *If $\tilde{y} \in \mathbb{R}^{\mathcal{A}}$ satisfies $\phi_\lambda(\tilde{y}) < 0$ for some $\lambda \in \Lambda_X(\mathcal{A})$, then $\exp \tilde{y} \notin C_X(\mathcal{A})^*$.*

Proof. We will find a vector $z \in \mathbb{R}^{\mathcal{A}}$ where $0 \leq z^T \exp y$ for all $\exp y \in C_X(\mathcal{A})^*$, and yet $z^T \exp \tilde{y} < 0$. By continuity, the condition that $0 \leq z^T \exp y$ for all $\exp y \in C_X(\mathcal{A})^*$ will imply the slightly stronger statement that $0 \leq z^T v$ for all $v \in C_X(\mathcal{A})^*$. Therefore z will evidently serve as a separating hyperplane to prove the desired claim. Let $\beta := \lambda^-$.

Since $\lambda \in \Lambda_X(\mathcal{A})$, Theorem 5.4 says that $\phi_\lambda(y) \geq 0$ whenever $\exp y \in C_X(\mathcal{A})^*$. Combine $\phi_\lambda(\tilde{y}) < 0$ with strict monotonicity of the exponential function to conclude

$$\exp(\phi_\lambda(\tilde{y})) < 1 \leq \exp(\phi_\lambda(y)) \quad \text{for all } \exp y \in C_X(\mathcal{A})^*. \quad (21)$$

Notice that taking a difference $\phi_\lambda(y) - \phi_\lambda(\tilde{y}) = \lambda_{\setminus\beta}^T(y_{\setminus\beta} + \tilde{y}_{\setminus\beta}) - y_\beta + \tilde{y}_\beta$ eliminates the support function term appearing in ϕ_λ . Defining $u = \phi_\lambda(\tilde{y})$, we multiply both sides of the non-strict inequality in (21) by $\exp(-u - \tilde{y}_\beta + y_\beta)$ to obtain

$$0 \leq \exp\left(\lambda_{\setminus\beta}^T(y_{\setminus\beta} - \tilde{y}_{\setminus\beta})\right) - \exp(-u - \tilde{y}_\beta + y_\beta). \quad (22)$$

Convexity of the exponential function tells us that $\exp\left(\lambda_{\setminus\beta}^T(y_{\setminus\beta} - \tilde{y}_{\setminus\beta})\right) \leq \lambda_{\setminus\beta}^T \exp(y_{\setminus\beta} - \tilde{y}_{\setminus\beta})$, where the right-hand-side may be rewritten using the Hadamard product

$$\lambda_{\setminus\beta}^T \exp(y_{\setminus\beta} - \tilde{y}_{\setminus\beta}) = \left(\lambda_{\setminus\beta} \circ \exp(-\tilde{y}_{\setminus\beta})\right)^T \exp(y_{\setminus\beta}).$$

Applying these observations to (22) gives

$$0 \leq \left(\lambda_{\setminus\beta} \circ \exp(-\tilde{y}_{\setminus\beta})\right)^T \exp(y_{\setminus\beta}) - (\exp(-u - \tilde{y}_\beta)) \exp(y_\beta). \quad (23)$$

Inequality (23) is essentially what we need to prove the lemma. Defining $z \in \mathbb{R}^{\mathcal{A}}$ by $z_\alpha = \lambda_\alpha \exp(-\tilde{y}_\alpha)$ for $\alpha \neq \beta$ and $z_\beta = -\exp(-u - \tilde{y}_\beta)$, we have that $0 \leq z^T \exp y$ for all $\exp y \in C_X(\mathcal{A})^*$. As explained at the beginning of this proof, we appeal to continuity to establish $0 \leq z^T v$ for all $v \in C_X(\mathcal{A})^*$. One may use $\lambda_{\setminus\beta}^T 1 = 1$ to trivially evaluate $z^T \exp(\tilde{y}) = 1 - \exp(-u)$, and since $u < 0$ by assumption on \tilde{y} , we conclude $z^T \exp(\tilde{y}) < 0$. \square

Proof of Theorem 5.6. By Theorem 5.4, we have the dual description $C_X(\mathcal{A})^* = \text{cl}\{\exp y : (y, 1) \in G_X(\mathcal{A})^*\}$. Applying Theorem 5.9 then gives

$$C_X(\mathcal{A})^* = \text{cl}\{\exp y : \phi_\lambda(y) \geq 0 \ \forall \lambda \in \Lambda_X^*(\mathcal{A})\}.$$

We rewrite the condition on $\phi_\lambda(y)$ as a condition on $v = \exp y$ using the power-cone formulation in Proposition 4.7. Since X is polyhedral, Theorem 3.7 tells us there are finitely many normalized X -circuits $\Lambda_X(\mathcal{A})$. We may therefore express $C_X(\mathcal{A})^*$ as a finite intersection of dual λ -witnessed AGE cones,

$$C_X(\mathcal{A})^* = \bigcap_{\lambda \in \Lambda_X^*(\mathcal{A})} C_X(\mathcal{A}, \lambda)^*.$$

Moreover, each dual λ -witnessed AGE cone $C_X(\mathcal{A}, \lambda)^*$ is an outer-approximation of the full-dimensional moment cone $\text{co}\{\exp(\mathcal{A}^T x) : x \in X\}$, hence there exists a point v_0 in the interior

of the moment cone where $v_0 \in \text{int } C_X(\mathcal{A}, \lambda)^*$ for all $\lambda \in \Lambda_X^*(\mathcal{A})$. Therefore, by [37, Corollary 16.4.2] we have

$$C_X(\mathcal{A}) = (C_X(\mathcal{A})^*)^* = \sum_{\lambda \in \Lambda_X^*(\mathcal{A})} (C_X(\mathcal{A}, \lambda)^*)^* = \sum_{\lambda \in \Lambda_X^*(\mathcal{A})} C_X(\mathcal{A}, \lambda),$$

which establishes the first part of the theorem.

For the second part of the theorem, suppose Λ is a proper subset of $\Lambda_X^*(\mathcal{A})$. Consider the set $C = \sum_{\lambda \in \Lambda} C_X(\mathcal{A}, \lambda)$ and its dual $C^* = \bigcap \{C_X(\mathcal{A}, \lambda)^* : \lambda \in \Lambda\}$. Clearly, since $C \subset C_X(\mathcal{A})$ we have $C^* \supset C_X(\mathcal{A})^*$ – we will show that this containment is *strict*, i.e., $C^* \not\supseteq C_X(\mathcal{A})^*$. Once this is done, duality will tell us that $C \subsetneq C_X(\mathcal{A})$.

Since C is contained within the signomial nonnegativity cone we again have that C^* contains the moment cone and so by Lemma 5.8 we have $C^* = \text{cl}(C^* \cap \mathbb{R}_{++}^A)$. Work with C^* over the positive orthant using Proposition 4.7 to express it as $C^* = \text{cl}\{\exp y : y \in Y\}$ for $Y := \{y : \phi_\lambda(y) \geq 0 \ \forall \lambda \in \Lambda\}$. By Lemma 5.12 there exists an element $\tilde{y} \in Y$ for which some $\lambda \in \Lambda_X^*(\mathcal{A}) \setminus \Lambda$ satisfies $\phi_\lambda(\tilde{y}) < 0$. Apply Lemma 5.14 to this pair $(\phi_\lambda, \tilde{y})$ to see that $\exp \tilde{y}$ can be separated from the closed convex set $C_X(\mathcal{A})^*$. We have therefore found a point \tilde{y} where $\exp \tilde{y} \in C^*$ and yet $\exp \tilde{y}$ can be separated from $C_X(\mathcal{A})^*$, so we conclude $C^* \not\supseteq C_X(\mathcal{A})^*$. \square

Before concluding this section we would like to point out a more general way to frame our analysis. Given a pair $(\lambda, a) \in \mathbb{R}^m \times \mathbb{R}$ where λ sums to zero and has exactly one negative component $\lambda_i = -1$, we have a power cone constraint $v_i \leq \exp(a) \prod_{j \neq i} v_j^{\lambda_j}$ which may be rewritten to $1 \leq \exp(a)v^\lambda$. Given a *set* of such pairs $P \subset \mathbb{R}^m \times \mathbb{R}$, we obtain the convex set

$$F(P) = \left\{ v \in \mathbb{R}_+^m : 1 \leq \exp(a)v^\lambda \ \forall (\lambda, a) \in P \right\}.$$

We have effectively shown that if $K = \text{co}(P \cup \{(\mathbb{0}, 1)\})$ is pointed and $F(P)$ intersects the positive orthant, then the unique minimum $P^* \subset P$ for which $F(P^*) = F(P)$ can be read off from the extreme rays of the polyhedral cone K .

6. EXTREME RAYS OF HALF-LINE SAGE CONES

In the previous section, we showed that by appropriate appeals to convex duality, one may derive representations of $C_X(\mathcal{A})$ with little to no redundancy. Here we build upon those results to completely characterize the extreme rays of the X -SAGE cone for the univariate case $X = [0, \infty)$.

Proposition 6.1. *For $\alpha_1 < \dots < \alpha_m$, the extreme rays of $C_{[0, \infty)}(\{\alpha_1, \dots, \alpha_m\})$ are:*

- (1) $\mathbb{R}_+ \cdot \exp(\alpha_1 x)$,
- (2) $\mathbb{R}_+ \cdot \{\exp(\alpha_2 x) - \exp(\alpha_1 x)\}$,
- (3) $\mathbb{R}_+ \cdot \{c_{i+1} \exp(\alpha_{i+1} x) + c_i \exp(\alpha_i x) + c_{i-1} \exp(\alpha_{i-1} x) : 2 \leq i \leq m-1\}$ with

$$c_{i+1} > 0, \quad c_{i-1} > 0, \quad \text{and} \quad c_i = - \left(\frac{c_{i-1}}{\lambda_{i-1}} \right)^{\lambda_{i-1}} \left(\frac{c_{i+1}}{\lambda_{i+1}} \right)^{\lambda_{i+1}},$$

where

$$\lambda_{i+1} = \frac{\alpha_i - \alpha_{i-1}}{\alpha_{i+1} - \alpha_{i-1}}, \quad \lambda_{i-1} = \frac{\alpha_{i+1} - \alpha_i}{\alpha_{i+1} - \alpha_{i-1}}, \quad \text{and} \quad \frac{c_{i-1}}{c_{i+1}} \geq \frac{\lambda_{i-1}}{\lambda_{i+1}}.$$

Proof. Let $\mathcal{A} = \{\alpha_1, \dots, \alpha_m\}$. By Theorem 5.6, all edge generators of $C_{[0, \infty)}(\mathcal{A})$ are either monomials or λ -witnessed AGE functions where λ is a reduced $[0, \infty)$ -circuit. By Example 5.7, $\Lambda_{[0, \infty)}^*(\mathcal{A}) = \{\delta_2 - \delta_1\} \cup \Lambda_{\mathbb{R}}^*(\mathcal{A})$. Since $n = 1$, Proposition 3.4 says all circuits λ have $|\text{supp } \lambda| \leq 3$. We therefore divide the proof into considering cases of monomials, and X -AGE functions with two or three terms.

First we address the monomials. Given $f(x) = \exp(\alpha_i x)$ with $i > 1$, we can write $f = f_1 + f_2$ with $f_1(x) = \exp(\alpha_i x) - \exp(\alpha_{i-1} x)$ and $f_2(x) = \exp(\alpha_{i-1} x)$ – the summand f_1 is nonnegative on $[0, \infty)$ because $\alpha_i > \alpha_{i-1}$, and f_2 is globally nonnegative. Therefore the only possible extremal monomial in $C_{[0, \infty)}(\mathcal{A})$ is $f(x) = \exp(\alpha_1 x)$. Since $X = [0, \infty)$, the leading term of any $g \in C_X(\mathcal{A})$ must have positive coefficient. Moreover, if g is not proportional to f , the leading term of g must have exponent greater than α_1 . Therefore any convex combination of AGE functions $g \in C_{[0, \infty)}(\mathcal{A})$ which are not proportional to f must disagree with $f(x)$ in the limit as x tends to infinity. We conclude f is extremal in $C_{[0, \infty)}(\mathcal{A})$.

Now we consider the 2-term case, where, by Example 5.7, we have to consider signomials of the form $f(x) = c_2 \exp(\alpha_2 x) - c_1 \exp(\alpha_1 x)$. We observe that f is nonnegative on $[0, \infty)$ if and only if $c_2 \geq c_1 \geq 0$, and furthermore that such signomials are nonextremal unless $c_1 = c_2$. To see that $f(x) = \exp(\alpha_2 x) - \exp(\alpha_1 x)$ is indeed extremal, note that f cannot be written as a convex combination involving any 3-term AGE functions, because any conic combination of 3-term AGE functions has a leading term with positive coefficient on $\exp(\alpha_i x)$ for some $i \geq 3$.

We have already proven cases (1) and (2) of the proposition. Using Example 5.7, we know that any extremal 3-term X -AGE function belongs to a λ -witnessed AGE cone where λ is a reduced \mathbb{R} -circuit. These reduced \mathbb{R} -circuits have the property $\text{supp } \lambda = \{i-1, i, i+1\}$ $\alpha_{i-1} \lambda_{i-1} + \alpha_{i+1} \lambda_{i+1} = \alpha_i$, $\lambda_i = -1$. Any X -AGE function with such a witness is nonnegative on all of \mathbb{R} . Therefore any 3-term X -AGE function f that is *extremal* in $C_{[0, \infty)}(\mathcal{A})$ is also extremal in $C_{\mathbb{R}}(\mathcal{A}) \subset C_{[0, \infty)}(\mathcal{A})$, which (by [17, Prop. 4.4]) implies

$$f(x) = c_{i+1} \exp(\alpha_{i+1} x) - \left(\left[\frac{c_{i+1}}{\lambda_{i+1}} \right]^{\lambda_{i+1}} \left[\frac{c_{i-1}}{\lambda_{i-1}} \right]^{\lambda_{i-1}} \right) \exp(\alpha_i x) + c_{i-1} \exp(\alpha_{i-1} x). \quad (24)$$

We have arrived at the final phase of proving part (3) of this proposition. By the equality case in the AM/GM inequality and using $\exp(\alpha_i x) = \left(\exp(\alpha_{i+1} x)^{\lambda_{i+1}} \right) \left(\exp(\alpha_{i-1} x)^{\lambda_{i-1}} \right)$, one finds the unique minimizer x^* for functions (24) satisfies

$$\left[\frac{c_{i+1} \exp(\alpha_{i+1} x^*)}{\lambda_{i+1}} \right] = \left[\frac{c_{i-1} \exp(\alpha_{i-1} x^*)}{\lambda_{i-1}} \right] \Leftrightarrow x^* = \ln \left(\frac{c_{i-1} \lambda_{i+1}}{c_{i+1} \lambda_{i-1}} \right) / (\alpha_{i+1} - \alpha_{i-1}).$$

If $V_i(\lambda, c) := (c_{i-1} \lambda_{i+1}) / (c_{i+1} \lambda_{i-1})$ satisfies $V_i(\lambda, c) < 1$, then $x^* < 0$ and by continuity we have $\inf\{f(x) : x \geq 0\} > 0$ – hence the condition $V_i(\lambda, c) \geq 1$ is necessary for extremality. Furthermore, if $V_i(\lambda, c) > 1$, then the unique minimizer of f given by (24) occurs at $x^* > 0$. Such f cannot be decomposed as a convex combination which involves 1-term or 2-term AGE functions (which have $f(x) > 0$ for $x > 0$), and cannot be written as a convex combination consisting solely of 3-term AGE functions [17, Proposition 4.4], therefore any f given by (24) with $V_i(\lambda, c) > 1$ is extremal in $C_{[0, \infty)}(\mathcal{A})$. All that remains is to show extremality of functions (24) with $V_i(\lambda, c) = 1$. This follows from the same argument as $V_i(\lambda, c) > 1$, but we must use

the stationarity condition $f'(0) = 0$ to preclude using 2-term extremal AGE functions in a decomposition of f . \square

7. DISCUSSION AND CONCLUSION

In this article we have introduced a convex-geometric notion of an X -circuit, which mediates a relationship between point sets $\mathcal{A} \subset \mathbb{R}^n$ and convex sets $X \subset \mathbb{R}^n$. By showing that this notion of an X -circuit allows an alternative construction of X -SAGE cones (Theorems 4.4 and 5.5) which cannot be relaxed (Theorem 5.6), we have demonstrated that conditional SAGE cones exhibit a substantially richer theory than ordinary SAGE cones. An essential property of this theory is that for general sets X it is not possible to recover an X -circuit $\lambda \in \Lambda_X(\mathcal{A}, \beta)$ given only information on the signs of its components. As a consequence of this last point – it is not possible to arrive at the concept of conditional SAGE certificates while relying on a “circuit number” approach using only the support of a given polynomial or signomial.

Two lines of theoretical investigations stand out for future work. First, there is the task of formally situating X -circuits in the context of matroid theory (in the case when X is a polyhedron). Here one can use an interpretation from Theorem 3.7, that X -circuits $\lambda \in \Lambda_X(\mathcal{A}, \beta)$ are outer normal vectors to facets of $-\mathcal{A}^T X + N_\beta^\circ$. A broader area of follow-up work is in-depth analysis of multiplicatively-convex sets $S \subset \mathbb{R}_+^m$ for which $\log(S) = \{t : \exp t \in S\}$ is convex. Some properties of this class of sets include closure under intersection, and closure under the induced-cone operation.

It is of interest to explore the use of the cones $C_X(\mathcal{A}, \lambda)$ when λ is not an X -circuit. Given a signomial $\sum_{\alpha \in \mathcal{A}} c_\alpha e^\alpha$ with numerical X -SAGE certificate $\{(c^{(\beta)}, \nu^{(\beta)})\}_{\beta \in \mathcal{A}}$, $c^{(\beta)} \in C_X(\mathcal{A}, \beta)$, $c \approx \sum_{\beta \in \mathcal{A}} c^{(\beta)}$, one could refine this certificate to higher precision by solving the power-cone program to decompose c as a sum of vectors in $C_X(\mathcal{A}, \lambda^{(\beta)})$ for $\lambda^{(\beta)} = \nu^{(\beta)} / |\nu^{(\beta)}|$. This would be helpful for large scale problems where $\{(c^{(\beta)}, \nu^{(\beta)})\}_{\beta \in \mathcal{A}}$ is computed with a first-order solver, or when X is an especially complicated spectrahedron. In the latter case, the standard description of $C_X(\mathcal{A})$ would be a mixed semidefinite and relative entropy program, while the formulations for $C_X(\mathcal{A}, \lambda^{(\beta)})$ would be pure power cone programs.

The two obstacles to using Theorem 5.6 in computation are that $|\Lambda_X^*(\mathcal{A})|$ can be exponential in $|\mathcal{A}|$ even when $X = \mathbb{R}^n$, and that finding X -circuits requires a procedure to identify extreme rays of a polyhedral cone. It is not known how severe this first problem is in practice. For the second problem one could focus on X -SAGE polynomials where $X = [-1, 1]^n$ or $X = [0, 1]^n$. The cones of such polynomials on $\mathcal{A} \subset \mathbb{N}^n$ are represented by $C_Y(\mathcal{A})$ for $Y = \{y \in \mathbb{R}^n : y \leq \mathbf{0}\}$, and finding $\Lambda_Y^*(\mathcal{A})$ is made easier by the fact that Y is a cone. The main benefit of this approach for polynomials is the prospect of computing conditional SAGE decompositions in exact arithmetic, especially for sparse polynomials of high degree.

We conclude by noting that although the “pure” conditional SAGE methodology is used only for convex constraint sets, additional nonconvex constraints can be accommodated with algebraic techniques. This can partly be seen in the original work of Chandrasekaran and Shah [4] and more so in the recent work [9].

REFERENCES

- [1] A. Agrawal, S. Diamond, and S. Boyd. Disciplined geometric programming. *Optim. Letters*, 13(5):961–976, 2019.
- [2] G. Averkov. Optimal size of linear matrix inequalities in semidefinite approaches to polynomial optimization. *SIAM J. Appl. Algebra and Geometry*, 3(1):128–151, 2019.
- [3] A. Ben-Tal and A. Nemirovski. *Lectures on Modern Convex Optimization*. SIAM, Philadelphia, 2001.
- [4] V. Chandrasekaran and P. Shah. Relative entropy relaxations for signomial optimization. *SIAM J. Optim.*, 26(2):1147–1173, 2016.
- [5] J. Dahl and E. Andersen. A primal-dual interior-point algorithm for nonsymmetric exponential-cone optimization. *Math. Program.*, 2021.
- [6] P. Dickinson and J. Povh. On an extension of Pólya’s Positivstellensatz. *J. Global Optimization*, 61:615–625, 2015.
- [7] M. Dressler, S. Ilman, and T. de Wolff. A Positivstellensatz for sums of nonnegative circuit polynomials. *SIAM J. Appl. Algebra and Geometry*, 1(1):536–555, 2017.
- [8] M. Dressler, A. Kurpisz, and T. de Wolff. Optimization over the Boolean hypercube via sums of nonnegative circuit polynomials. In *43rd International Symposium on Mathematical Foundations of Computer Science*, volume 117 of *LIPICs. Leibniz Int. Proc. Inform.*, pages 82:1–82:17. Schloss Dagstuhl, 2018.
- [9] M. Dressler and R. Murray. Algebraic perspectives on signomial optimization. Preprint, ArXiv:2107.00345, 2021.
- [10] A. A Ergür, G. Paouris, and J. M. Rojas. Tropical varieties for exponential sums. *Math. Annalen*, 377:863–882, 2020.
- [11] F. Fillastre and I. Izmistiev. Shapes of polyhedra, mixed volumes and hyperbolic geometry. *Mathematika*, 63(1):124–183, 2017.
- [12] J. Forsgård and T. de Wolff. The algebraic boundary of the SONC cone. Preprint, arXiv:1905.04776, 2019.
- [13] I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky. *Discriminants, Resultants, and Multidimensional Determinants*. Springer, 1994.
- [14] S. Ilman and T. de Wolff. Amoebas, nonnegative polynomials and sums of squares supported on circuits. *Res. Math. Sci.*, 3(paper no. 9), 2016.
- [15] W. Jarczyk and J. Matkowski. On Mulholland’s inequality. *Proc. Amer. Math. Soc.*, 130(11):3243–3247, 2002.
- [16] O. Karaca, G. Darivianakis, P. Beuchat, A. Georghiou, and J. Lygeros. The REPOP toolbox: Tackling polynomial optimization using relative entropy relaxations. In *20th IFAC World Congress, IFAC PapersOnLine*, volume 50(1), pages 11652–11657. Elsevier, 2017.
- [17] L. Katthän, H. Naumann, and T. Theobald. A unified framework of SAGE and SONC polynomials and its duality theory. *Math. Computation*, 90:1297–1322, 2021.
- [18] J.B. Lasserre. *Moments, Positive Polynomials and their Applications*. Imperial College Press, London, 2010.
- [19] S. Müller, E. Feliu, G. Regensburger, C. Conradi, A. Shiu, and A. Dickenstein. Sign conditions for injectivity of generalized polynomial maps with applications to chemical reaction networks and real algebraic geometry. *Found. Comp. Math.*, 16(1):69–97, 2015.
- [20] S. Müller, J. Hofbauer, and G. Regensburger. On the bijectivity of families of exponential/generalized polynomial maps. *SIAM J. Appl. Algebra and Geometry*, 3(3):412–438, 2019.
- [21] R. Murray. Sageopt 0.5.3, 2020. DOI:10.5281/ZENODO.4017991.
- [22] R. Murray. *Applications of Convex Analysis to Signomial and Polynomial Nonnegativity Problems*. PhD thesis, California Institute of Technology, 6 2021.
- [23] R. Murray, V. Chandrasekaran, and A. Wierman. Newton polytopes and relative entropy optimization. *Found. Comput. Math.*, 21:1703–1737, 2021.
- [24] R. Murray, V. Chandrasekaran, and A. Wierman. Signomial and polynomial optimization via relative entropy and partial dualization. *Math. Program. Comput.*, 13:257–295, 2021.

- [25] K. G. Murty and S. N. Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Math. Program.*, 39(2):117–129, 1987.
- [26] H. Naumann and T. Theobald. The \mathcal{S} -cone and a primal-dual view on second-order representability. *Beiträge Algebra Geom.*, 62:229–249, 2021.
- [27] H. Naumann and T. Theobald. Sublinear circuits for polyhedral sets. *Vietnam J. Math.*, 2022.
- [28] C. P. Niculescu. Convexity according to the geometric mean. *Math. Inequalities & Applications*, 3:155–167, 2000.
- [29] C. Nowzari, V. M. Preciado, and G. J. Pappas. Optimal resource allocation for control of networked epidemic models. *IEEE Transactions on Control of Network Systems*, 4(2):159–169, 2017.
- [30] M. Özdemir, Ç. Yildiz, and M. Gürbüz. A note on geometrically convex functions. *J. Inequalities and Applications*, 2014(1):180, 2014.
- [31] B. Öztürk and A. Saab. Optimal aircraft design decisions under uncertainty via robust signomial programming. In *AIAA Aviation 2019 Forum*. American Institute of Aeronautics and Astronautics, article no. 2019–3351, 2019.
- [32] C. Pantea, H. Koepl, and G. Craciun. Global injectivity and multiple equilibria in uni- and bi-molecular reaction networks. *Discrete and Continuous Dynamical Systems - Series B*, 17(6):2153–2170, 2012.
- [33] D. Papp. Duality of sum of nonnegative circuit polynomials and optimal SONC bounds. Preprint, arXiv:1912.04718, 2019.
- [34] V. Powers and T. Wörmann. An algorithm for sums of squares of real polynomials. *J. Pure and Applied Algebra*, 127(1):99–104, 1998.
- [35] V. M. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. J. Pappas. Optimal resource allocation for network protection against spreading processes. *IEEE Transactions on Control of Network Systems*, 1(1):99–108, 2014.
- [36] B. Reznick. Forms derived from the arithmetic-geometric inequality. *Math. Annalen*, 283(3):431–464, 1989.
- [37] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1997.
- [38] B. Sturmfels. *Gröbner Bases and Convex Polytopes*. Amer. Math. Soc., 1996.
- [39] A. H. Wang, P. Jaini, Y. Yu, and P. Poupert. A Positivstellensatz for conditional SAGE signomials. Preprint, arXiv:2003.03731, 2020.
- [40] J. Wang. Nonnegative polynomials and circuit polynomials. *SIAM J. Appl. Algebra and Geometry*, 2022.
- [41] J. Wang and V. Magron. A second order cone characterization for sums of nonnegative circuits. In *Proc. International Symposium on Symbolic and Algebraic Computation (ISSAC)*, pages 450–457. ACM, 2020.
- [42] M. York, W. Hoburg, and M. Drela. Turbofan engine sizing and tradeoff analysis via signomial programming. *Journal of Aircraft*, 55(3):988–1003, 2018.
- [43] G. M. Ziegler. *Lectures on Polytopes*. Springer, New York, 1995.

8. APPENDIX

8.1. Propositions regarding convex analysis. The following pair of propositions are used in the proofs of Theorem 4.4 and Lemma 5.11.

Proposition 8.1. *For fixed λ in the interior of the m -dimensional probability simplex and $c = (c_0, c_1, \dots, c_m) \in \mathbb{R}^{m+1}$ with $(c_1, \dots, c_m) \geq \mathbf{0}$, we have*

$$-c_0 \leq \prod_{i=1}^m [c_i/\lambda_i]^{\lambda_i} \quad \Leftrightarrow \quad \text{some } \nu \in \mathbb{R}_+^m \text{ satisfies } \nu \parallel \lambda \text{ and } D(\nu, c_{\setminus 0}) - \mathbf{1}^T \nu \leq c_0$$

– where $\nu \parallel \lambda$ means ν is proportional to λ .

Proof. The claim is trivial when $c_0 \geq 0$, and so we consider $c_0 < 0$. Note that in this case, $\prod_{i=1}^m [c_i/\lambda_i]^{\lambda_i}$ must be positive, and $D(\nu, c_{\setminus 0})$ must be finite: both of these conditions occur

precisely when $c_i > 0$ for all $1 \leq i \leq m$. We therefore can rewrite $-c_0 = |c_0| \leq \prod_{i=1}^m [c_i/\lambda_i]^{\lambda_i}$ as $1 \leq \prod_{i=1}^m [c_i/(|c_0|\lambda_i)]^{\lambda_i}$, and by taking the log of both sides, obtain $D(\nu, c_{\setminus 0}) - \mathbb{1}^T \nu \leq c_0$ for $\nu = |c_0|\lambda$. For the other direction, one may write the proportionality relationship $\nu \parallel \lambda$ as $\nu = s\lambda$, and minimize $D(s\lambda, c_{\setminus 0}) - s$ over $s \geq 0$ to obtain $-\prod_{i=1}^m [c_i/\lambda_i]^{\lambda_i}$. \square

Proposition 8.2. *Suppose $S \subset \mathbb{R}^m \setminus \{\mathbf{0}\}$ is compact (not necessarily convex) and set $T = \text{co } S$. If it is known a-priori that $\text{cl } T$ contains no lines, then $T = \text{cl } T$ is closed.*

Proof. Since $\text{cl } T$ is pointed, there exists a distinguished element $t^* \in T$ for which $(t^*)^T t > 0$ for all $t \in (\text{cl } T) \setminus \{\mathbf{0}\}$. Consider the set $H = \{t \in T : (t^*)^T t = 1\}$ – it is clear that H is bounded, $\text{co } H = T$, and $0 \notin H$. If H is closed, then by [37, Corollary 9.6.1] we will have that $\text{co } H = T$ is also closed. We show that H is closed by directly considering sequences in H . We express these sequences with the help of the m -fold Cartesian product $S^m = S \times \dots \times S$.

Let $(h^{(k)})_{k \in \mathbb{N}} \subset H$ have a limit in \mathbb{R}^m . Since H is of dimension at most $m-1$ and is generated by S , Carathéodory's Theorem tells us that there exists a vector $\lambda^{(k)} \in \mathbb{R}_+^m$ and a block vector $q^{(k)} = (s_1^{(k)}, \dots, s_m^{(k)}) \in S^m$ where

$$h^{(k)} = \sum_{i=1}^m \lambda_i^{(k)} s_i^{(k)}.$$

Since S is compact, the continuous function $s \mapsto (t^*)^T s$ attains a minimum on $s^* \in S$ – since S does not contain zero, we have that $(t^*)^T (s^*) = a > 0$. It follows that each $\lambda_i^{(k)}$ appearing in the expression for $h^{(k)}$ is bounded above by $1/a < \infty$. The sequences $(\lambda^{(k)})_{k \in \mathbb{N}} \subset [0, 1/a]^m$ and $(q^{(k)})_{k \in \mathbb{N}} \subset S^m$ are bounded, and therefore $((\lambda^{(k)}, q^{(k)}))_{k \in \mathbb{N}}$ has a convergent subsequence. The limits $\lambda^{(\infty)}$ and $q^{(\infty)}$ of these convergent subsequences must belong to $[0, 1/a]^m$ and S^m , respectively. By continuity, we have

$$h^{(\infty)} := \lim_{k \rightarrow \infty} h^{(k)} = \sum_{i=1}^m \lambda_i^{(\infty)} s_i^{(\infty)},$$

hence $h^{(\infty)} \in H$. Since we have shown that all convergent sequences in H converge to a *point in* H , we have that H is closed. \square

Our next proposition is provided for the reader's convenience.

Proposition 8.3. *Let $X \subset \mathbb{R}^n$ be a convex cone and consider a matrix A in $\mathbb{R}^{n \times m}$. We have $(A^T X)^* = \ker A + A^\dagger X^*$, where $A^\dagger \in \mathbb{R}^{m \times n}$ is the Moore-Penrose pseudo-inverse of A .*

Proof. Because $A^T X$ is contained in the subspace $\text{range } A^T$, its dual cone is invariant under translation by vectors in the orthogonal complement $(\text{range } A^T)^\perp = \ker A$. In particular, $(A^T X)^* = \ker A + K$ for a convex cone $K \subset \text{range } A^T$. We need to show that $K = A^\dagger X^*$.

The definition of the Moore-Penrose pseudo-inverse ensures that $y \in \text{range } A^T$ holds if and only if $A^\dagger A y = y$. We can therefore compute K as follows

$$\begin{aligned} K &= \{y \in \mathbb{R}^m : y^T z \geq 0 \forall z \in A^T X\} \cap \{y : A^\dagger A y = y\} \\ &= \{A^\dagger A y : (A y)^T x \geq 0 \forall x \in X\} \cap \{y : A^\dagger A y = y\} \\ &= \{A^\dagger w : w \in \mathbb{R}^n, w^T x \geq 0 \forall x \in X\} \cap (\text{range } A^T) \\ &= \{A^\dagger w : w \in X^*\}. \end{aligned}$$

The transitions from line to line are as follows. First, substitute $A^\dagger Ay$ for y , express $z = A^T x$ for some $x \in X$, and rewrite $y^T(A^T x) = (Ay)^T x$. Then, substitute $w := Ay$ and simplify the expression for the range of A^T . Finally, apply the definition of the dual cone X^* and use the pseudo-inverse identity $A^\dagger AA^\dagger x = A^\dagger x$ for all $x \in \mathbb{R}^n$. \square

8.2. Definitions from convex analysis. A *face* of a convex set $S \subset \mathbb{R}^n$ is any closed convex $F \subset S$ with the following property: if the line segment $[s_1, s_2] := \{\lambda s_1 + (1-\lambda)s_2 : 0 \leq \lambda \leq 1\}$ is contained in S and the relative interior of $[s_1, s_2]$ hits F , then the entirety of $[s_1, s_2]$ is contained in F . The *dimension* $\dim S$ of a convex set S is the dimension of the smallest affine space containing S . Every nonempty convex set S has a nonempty *relative interior* $\text{ri } S$, which is the interior of S under the topology induced by its affine hull. A set $K \subset \mathbb{R}^n$ is called a *cone* if it is closed under dilation: $\{\lambda x : x \in K\} \subset K$ for all $\lambda > 0$. The *extreme rays* of a pointed convex cone K are its faces of dimension one. To any convex cone K we associate the *dual cone* $K^* := \{y : y^T x \geq 0 \forall x \in K\}$ and the *polar* $K^\circ = -K^*$. The *conic hull* of a set S , denoted $\text{co } S$, is the set formed by adjoining the origin to the smallest convex cone containing S .

RILEY MURRAY: UNIVERSITY OF CALIFORNIA, BERKELEY - DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCES (rjmurray@berkeley.edu). PREVIOUSLY, CALIFORNIA INSTITUTE OF TECHNOLOGY - DEPARTMENT OF COMPUTING AND MATHEMATICAL SCIENCES

HELEN NAUMANN, THORSTEN THEOBALD: GOETHE-UNIVERSITÄT, FB 12 – INSTITUT FÜR MATHEMATIK, POSTFACH 11 19 32, D-60054 FRANKFURT AM MAIN, GERMANY ({naumann,theobald}@math.uni-frankfurt.de)