

**Max-Planck-Institut
für Mathematik
in den Naturwissenschaften
Leipzig**

**The X-circuits Behind Conditional
SAGE Certificates**

by

Riley Murray, Helen Naumann, and Thorsten Theobald

Preprint no.: 65

2020



THE X -CIRCUITS BEHIND CONDITIONAL SAGE CERTIFICATES

RILEY MURRAY, HELEN NAUMANN, AND THORSTEN THEOBALD

ABSTRACT. Conditional SAGE certificates are a decomposition method to prove nonnegativity of a signomial or polynomial over some subset X of Euclidean real space. In the case when X is convex, membership in the signomial “ X -SAGE cone” can be completely characterized by a relative entropy program involving the support function of X . Following promising computational experiments, and a recently proven completeness result for a hierarchy of X -SAGE relaxations for signomial optimization, we undertake a structural analysis of signomial X -SAGE cones.

Our approach begins by determining a suitable notion of an “ X -circuit,” in such a way as to generalize classical affine-linear simplicial circuits from matroid theory. Our definition of an X -circuit is purely convex-geometric, with no reference to signomials or SAGE certificates. We proceed by using X -circuits to characterize the more elementary “ X -AGE cones” which comprise a given X -SAGE cone. Our deepest results are driven by a duality theory for X -circuits, which is applicable to primal and dual X -SAGE cones in their usual forms, as well as to a certain logarithmic transform of the dual cone. In conjunction with a notion of reduced X -circuits this facilitates to characterize the extreme rays of the X -SAGE cones. Our results require no regularity conditions on X beyond those which ensure a given X -SAGE cone is proper; particularly strong conclusions are obtained when X is a polyhedron.

1. INTRODUCTION

A circuit of an affine-linear matroid over \mathbb{R} is a subset A of a finite ground set $\mathcal{A} \subset \mathbb{R}^n$, where A is affinely dependent but any proper subset of A is affinely independent; simplicial circuits are those A with k elements whose convex hulls have $k - 1$ extreme points. Using $\mathbb{R}^{\mathcal{A}}$ to denote the set of vectors with components indexed by $\alpha \in \mathcal{A}$, simplicial circuits are representable by sparse vectors $\lambda \in \mathbb{R}^{\mathcal{A}}$ giving the barycentric coordinates of a circuit’s nonextremal element.

The link between simplicial circuits, convexity theory and algebraic optimization has manifested in certain certificates of nonnegativity for signomials (functions $x \mapsto \sum_{\alpha \in \mathcal{A}} c_{\alpha} \exp(\alpha^T x)$) and polynomials. The earliest results here are due to Reznick [1], with a recent resurgence marked by the works of Pantea, Koeppl, and Craciun [2], Iliman and de Wolff [3], and Chandrasekaran and Shah [4]. Both Pantea et al. and Iliman and de Wolff recognized that given a polynomial $f(x) = \sum_{\alpha \in \mathcal{A}} c_{\alpha} \prod_{i=1}^n x_i^{\alpha_i}$ whose support $\text{supp}(c) := \{\alpha : c_{\alpha} \neq 0\}$ is a simplicial circuit, there exist necessary and sufficient tests for \mathbb{R}_+^n - or \mathbb{R}^n -nonnegativity involving c and the circuit vector “ λ ” corresponding to $\text{supp}(c)$. Chandrasekaran and Shah showed if a signomial

Date: June 15, 2020.

2010 Mathematics Subject Classification. 14P05, 52A20, 90C30.

Key words and phrases. Sums of arithmetic-geometric exponentials, positive signomials, signomial optimization, exponential sums, sums of nonnegative circuit polynomials (SONC), positive polynomials, convex analysis, multiplicative convexity, log convex sets, relative entropy, power cones.

$f(x) = \sum_{\alpha \in \mathcal{A}} c_\alpha \exp(\alpha^T x)$ has at most one negative coefficient, nonnegativity of f is equivalent to the existence of a $\nu \in \mathbb{R}^A$ which certifies a particular relative entropy inequality over (ν, c) ; such nonnegative signomials were termed *AM/GM-exponentials* or *AGE functions*. Although circuits do not manifest explicitly in the theory of AGE functions, it has been shown that all functions which generate an extreme ray of an “AGE cone” either have $|\text{supp}(c)| = 1$ or $\text{supp}(c)$ is a simplicial circuit [5].

The terms *SONC* and *SAGE* traditionally referred to methods for certifying polynomial and signomial nonnegativity based on decomposition into sums of nonnegative circuit polynomials¹ and sums of AGE functions respectively. However, in view of results by Wang concerning SONC’s ability to certify nonnegativity of “AGE-like” polynomials [6], and Murray, Chandrasekaran, and Wierman’s proof that the cone of SONC polynomials can be represented by a projection of the cone of SAGE signomials [5], it is appropriate to consider SONC and SAGE as equivalent to one another for purposes of certain structural analysis. Although SAGE is extremely efficient from a computational complexity perspective, the circuit approach is generally more fruitful in answering structural questions. For example, the “reduced circuits” of Katthän, Naumann, and Theobald were used to completely characterize the extreme rays of the SAGE cones [7]. Subsequently, Forsgård and de Wolff employed the theories of regular subdivisions, A -discriminants and tropical geometry to study (among other things) how circuits affect the algebraic boundary of the signomial SAGE cone [8].

The SAGE and SONC nonnegativity certificates were originally developed as a means to certify *global* nonnegativity. When these certificates have been used in service of certifying nonnegativity over some $X \subsetneq \mathbb{R}^n$ (as is done to produce convex relaxations to constrained optimization problems), the tendency has been to adopt a representation of $X = \{x : g(x) \geq 0\}$ and subsequently appeal to the minimax inequality [4, 9, 10, 11, 12].

In a recent and major step forward, Murray, Chandrasekaran and Wierman extended SAGE certificates to constrained settings with convex feasible sets X in a way that does not rely on the minimax inequality [13]. It rests upon the key observation that the nonnegativity of a signomial $f(x) = \sum_{\alpha \in \mathcal{A}} c_\alpha \exp(\alpha^T x)$ on the convex set X with at most one negative coefficient c_β can be formulated exactly as a relative entropy program employing the support function of X . Murray et al. called their method *conditional SAGE*, out of a dual point of view which connects the method to moment problems in conditional probability. For brevity and concreteness, the term “*X-SAGE*” is often used in place of “conditional SAGE.”

The usefulness of conditional SAGE certificates was initially demonstrated by computational experiments [13]. Very recently, A. Wang et al. have developed a uniquely general convergence result for an *X-SAGE* hierarchy of convex relaxations to signomial optimization problems [14]. The primary motivation for the current article is to advance the understanding of the structure of signomial *X-SAGE* cones, and to establish a versatile access to the *X-SAGE* cones in terms of circuits. As we show, a suitable *sublinear* generalization of affine-linear circuits to the constrained setting provides an elegant convex-combinatorial framework with deep connections to polyhedral convexity. This notion’s resulting interaction with duality and multiplicative-convexity enables the characterization of *X-SAGE* cones in a surprisingly comprehensive way,

¹*nonnegative circuit polynomial* being shorthand for *nonnegative polynomial supported on a simplicial circuit*

even when X is neither polyhedral nor compact. These X -circuits exhibit a rich interplay of real algebraic geometry and optimization, as well as discrete and convex geometry.

The remainder of the article is organized as follows. In Section 2 we introduce assumptions used in the article, establish notation and definitions, and provide minimal background on ordinary and conditional SAGE certificates.

Section 3 defines our notion of an “ (\mathcal{A}, X) -circuit” (usually abbreviated to “ X -circuit”) as a vector ν^* satisfying a local strict-sublinearity condition for the augmented support function $\nu \mapsto \sup\{-\nu^T \mathcal{A}^T x : x \in X\}$. This builds upon and generalizes the classical notion of matroid-theoretic affine-linear “ \mathbb{R}^n -circuits.” The essential difference between X -circuits and \mathbb{R}^n -circuits, is that for some sets X , it is not possible to recover an X -circuit given only information on the signs of its components. The two main results of Section 3 (Theorems 3.6 and 3.7) are purely geometric statements on X -circuits with no reference to SAGE or signomials.

Section 4 concerns how the theory of X -circuits reveals structure in the “ X -AGE cones” that comprise an X -SAGE cone. Theorem 4.2 particularly shows that in order for an X -AGE function to be extremal in a given X -AGE cone, it is necessary that the auxiliary variable in its membership-certifying relative entropy inequality (to be described in Section 2.1) defines an X -circuit. This necessary condition for extremality leads us to define λ -witnessed AGE functions, which serve as a more basic building block for X -AGE cones. Propositions 4.9 and 4.10 provide power-cone representations for primal and dual cones of λ -witnessed AGE functions. While developing this section’s results for general convex sets X , we take small detours to derive additional results in the case when X is a polyhedron (Theorem 4.4, Corollaries 4.7 and 4.8).

Section 5 proves three major results on cones of X -SAGE signomials. At the heart of these results is the striking fact that whenever X is a convex set, a dual X -SAGE cone is not only convex in the classical sense, but also convex after a logarithmic transformation $S \mapsto \log S := \{t : \exp t \in S\}$. The property of a set being convex under this logarithmic transformation is known by various names, including *log convexity* [15], *geometric convexity* [16, 17], or *multiplicative convexity* [18],² and has previously been touched upon in the literature on ordinary SAGE certificates [5, 7, 19]. Theorem 5.3 provides an explicit description of the dual X -SAGE cone’s logarithmic transform, as an affine slice of the dual to a cone generated by X -circuits. We go on to prove Theorems 5.5 and 5.6, which concern representing an X -SAGE cone in terms of λ -witnessed AGE cones, where λ satisfies a reducedness condition which is far stronger than being an X -circuit (Definition 5.4). Theorem 5.6 specifically concerns the polyhedral case; its proof employs a technical lemma on separation theory of sets which are simultaneously classically-convex and multiplicatively-convex.

The development from Sections 3 through 5 is increasingly focused and goal-oriented. Section 6 takes a step back, and asks how our results manifest concretely, or how they may be understood without appeals to convex duality. In Theorem 6.1 we give a primal-only argument to demonstrate nonextremality of certain λ -witnessed AGE functions, where λ is a nonreduced X -circuit. Proposition 6.3 characterizes extreme rays for X -SAGE cones in the univariate case

²The class of functions corresponding to log-convex sets are those f satisfying $f(x^\theta y^{(1-\theta)}) \leq f(x)^\theta f(y)^{(1-\theta)}$ for $\theta \in (0, 1)$ and $x, y > 0$, or equivalently, those functions for which $\log x \mapsto \log f(x)$ is convex in the classical sense. Study of these “log-log convex” functions is far more common than study of log-convex sets themselves.

$X = [0, \infty)$; considering this specific case helps highlight distinctions between reduced and nonreduced X -circuits, and subtleties in being extremal in an X -SAGE cone versus being a λ -witnessed AGE function where λ is a reduced X -circuit.

We conclude the article with a discussion on a wide range of open problems. Among these open problems are questions regarding (\mathcal{A}, X) -circuits as purely geometric objects, the theory of multiplicatively-convex analysis, and the semidefinite representability of X -SAGE cones.

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.

2. PRELIMINARIES

All logarithms are base- e , where e is Euler's number. Throughout this article, $X \subset \mathbb{R}^n$ is closed, convex, and nonempty. The set $\mathcal{A} \subset \mathbb{R}^n$ is nonempty and finite; its elements are interpreted as linear functions over X . From \mathcal{A} we consider the set " $\mathbb{R}^{\mathcal{A}}$ " of vectors with components indexed by $\alpha \in \mathcal{A}$. Indexing according to elements of \mathcal{A} allows us to specify a signomial solely by a vector $c \in \mathbb{R}^{\mathcal{A}}$, via

$$f(x) = \sum_{\alpha \in \mathcal{A}} c_{\alpha} \exp(\alpha^T x).$$

The *support* of a signomial $f(x) = \sum_{\alpha \in \mathcal{A}} c_{\alpha} \exp(\alpha^T x)$ is $\{\alpha : c_{\alpha} \neq 0\}$. A *posynomial* is a signomial with $c \in \mathbb{R}_+^{\mathcal{A}}$. We often use the elementwise extension $\exp : \mathbb{R}^{\mathcal{A}} \rightarrow \mathbb{R}^{\mathcal{A}}$.

The set \mathcal{A} is routinely interpreted as a linear operator $\mathcal{A} : \mathbb{R}^{\mathcal{A}} \rightarrow \mathbb{R}^n$ by $\mathcal{A}\nu = \sum_{\alpha \in \mathcal{A}} \alpha \nu_{\alpha}$; the corresponding adjoint operator is denoted \mathcal{A}^T . We only consider data (\mathcal{A}, X) where the functions $\{x \mapsto \exp(\alpha^T x)\}_{\alpha \in \mathcal{A}}$ are linearly independent on X ; this assumption is necessary to prevent the nonnegativity cone from containing a lineality space. A direct consequence of this assumption is that the moment cone $\text{co}\{\exp(\mathcal{A}^T x) \in \mathbb{R}^{\mathcal{A}} : x \in X\}$ is full-dimensional.

2.1. The basics of X -SAGE signomials. A signomial with coefficient vector $c \in \mathbb{R}^{\mathcal{A}}$ is called X -AGE if it is nonnegative on X , and c contains at most one negative component. The cone of X -AGE signomials with support contained in \mathcal{A} and free term $\beta \in \mathcal{A}$ is denoted $C_X(\mathcal{A}, \beta)$. We often overload notation and consider $C_X(\mathcal{A}, \beta)$ as a cone of coefficient vectors in $\mathbb{R}^{\mathcal{A}}$.

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

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

It is possible to efficiently check membership in an X -SAGE cone whenever X is a tractable convex set. The technical sense in which X must be "tractable" is that we need a method to optimize over the epigraph of the set's *support function*

$$\sigma_X(y) = \sup_{x \in X} y^T x.$$

All variations of SAGE certificates prominently feature the *relative entropy function*

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

which is continuously extended to $\mathbb{R}_+^{\mathcal{A}} \times \mathbb{R}_+^{\mathcal{A}}$.

Throughout this article, vectors $\beta \in \mathcal{A}$ have associated sets $N_\beta = \{\nu \in \mathbb{R}_+^{\mathcal{A} \setminus \beta} \times \mathbb{R} : 1^T \nu = 0\}$, where $\mathcal{A} \setminus \beta$ abbreviates $\mathcal{A} \setminus \{\beta\}$. For $w \in \mathbb{R}^{\mathcal{A}}$ and $\beta \in \mathcal{A}$, $w_{\setminus \beta}$ is the vector in $\mathbb{R}^{\mathcal{A} \setminus \beta}$ obtained by dropping w_β from w .

Proposition 2.2 (Theorem 6 of [13]). *A signomial $f(x) = \sum_{\alpha \in \mathcal{A}} c_\alpha \exp(\alpha^T x)$ belongs to $C_X(\mathcal{A}, \beta)$ if and only if $c_\alpha \geq 0$ for all $\alpha \in \mathcal{A} \setminus \{\beta\}$ and some $\nu \in \mathbb{R}^{\mathcal{A}}$ satisfies*

$$\nu \in N_\beta \quad \text{and} \quad \sigma_X(-\mathcal{A}\nu) + D(\nu_{\setminus \beta}, e_{c_{\setminus \beta}}) \leq c_\beta. \quad (1)$$

2.2. The role of circuits in ordinary SAGE cones. The circuits used in the SONC or \mathbb{R}^n -SAGE literature are defined with respect to *affine* independence. Specifically, a circuit is a minimally-supported nonzero vector $\lambda \in \ker \mathcal{A} \subset \mathbb{R}^{\mathcal{A}}$ for which $\sum_{\alpha \in \mathcal{A}} \lambda_\alpha = 0$. The SONC or \mathbb{R}^n -SAGE literature primarily considers *simplicial circuits*, which have the additional property that λ contains a single negative component, say $\lambda_\beta < 0$. We usually canonically rescale $\lambda_\beta = -1$, so that simplicial circuits satisfy $\sum_{\alpha \in \text{supp } \lambda \setminus \{\beta\}} \alpha \lambda_\alpha = \beta(-\lambda_\beta) = \beta$. Simplicial circuits therefore certify that some “ β ” belongs to the relative interior of the point set $\text{supp } \lambda \setminus \{\beta\} \subset \mathcal{A}$.

Conditional SAGE certificates do not use circuits, but they do use vectors “ $\nu \in (\mathbb{R}_+^{\mathcal{A} \setminus \beta}) \times \mathbb{R}^\beta$ ” which sum to zero, and ν appears as an additive term $+\sigma_X(-\mathcal{A}\nu)$ in (1). Considering $X = \mathbb{R}^n$ specifically, we have that $\sigma_{\mathbb{R}^n}(-\mathcal{A}\nu) = 0$ when $\nu \in \ker \mathcal{A}$, and $\sigma_{\mathbb{R}^n}(-\mathcal{A}\nu) = +\infty$ otherwise. For an \mathbb{R}^n -AGE certificate to be nontrivial, we need that $\nu \in \ker \mathcal{A}$ contains exactly one negative component, and that $\sum_{\alpha \in \mathcal{A}} \nu_\alpha = 0$. Such a vector ν would define a circuit if it were a *minimally supported* element of $\ker \mathcal{A}$. Many AGE certificates (ν, c) do not satisfy the property of being minimally supported, and so AGE certificates do not universally involve circuits. Nevertheless, [5, Theorem 4] says that any AGE certificate (ν, c) can be mapped to a collection of AGE certificates $\{(\nu^{(i)}, c^{(i)})\}_{i=1}^\ell$ for which $\nu^{(i)}$ are minimally supported, and $\sum_{i=1}^\ell c^{(i)} = c$. In this way, circuits are not needed for computations with \mathbb{R}^n -SAGE certificates, but they do determine the extreme rays of $C_{\mathbb{R}^n}(\mathcal{A})$.

2.3. Convex analysis. We routinely cite results from convex analysis, as contained in Rockafellar’s [20]. The following terminology is generally chosen to match those of Rockafellar.

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 . We sometimes write $F \trianglelefteq S$ to indicate that F is a face of S . The *dimension* $\dim S$ of a convex set S is the dimension of the smallest affine space containing S ; *extreme points* of a convex set are its faces of dimension zero. The *relative interior* of a convex set, denoted $\text{ri } S$, is the interior of S under the topology induced by its affine hull. The *polar* of a convex set S is $S^\circ := \{t : \sigma_S(t) \leq 1\}$.

A set $K \subset \mathbb{R}^n$ is called a *cone* if it is closed under dilation $K \supset \{\lambda x : x \in K\}$ for all $\lambda > 0$. A convex cone $K \subset \mathbb{R}^n$ is *pointed* if it contains no lines. The *extreme rays* of a convex cone

K are its faces of dimension 1. 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 . To any convex cone K we associate the *dual cone* $K^* := \{y : y^T x \geq 0 \forall x \in K\} = -K^\circ$. 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 . Convex sets $S \subset \mathbb{R}^n$ have convex *induced cones* $\text{indco}(S) := \text{cl}\{(s, \mu) : \mu > 0, s/\mu \in S\} \subset \mathbb{R}^{n+1}$ and *recession cones* $\text{rec} S := \{t : \exists s \in S \text{ such that } s + \lambda t \in S \forall t \geq 0\}$.

We call a set a *polyhedron* if it can be represented by the intersection of finitely many half-spaces; *polytopes* are the bounded polyhedra.

3. X -CIRCUITS INDUCED BY A POINT-SET.

Recall from Section 2.1 the definition $N_\beta = \{\nu \in \mathbb{R}_+^{\mathcal{A} \setminus \beta} \times \mathbb{R} : 1^T \nu = 0\}$. This definition implies that every $\nu \in N_\beta$ has $\nu_\beta = -\sum_{\alpha \in \mathcal{A} \setminus \beta} \nu_\alpha \leq 0$, and $\|\nu\|_\infty = |\nu_\beta|$.

Definition 3.1. A vector $\nu^* \in N_\beta$ is an (\mathcal{A}, X) -*circuit* (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 $[\nu^{(1)}, \nu^{(2)}]$.

Remark 3.2. If no special properties are assumed for X , we often drop prefixes from terms “ (\mathcal{A}, X) -circuit” and “ X -circuit.” Exceptions to this rule are made when additional distinction from the classical case $X = \mathbb{R}^n$ is deemed worthwhile. Unqualified use of the word “circuit” does not default to $X = \mathbb{R}^n$.

There are some simple properties we wish to emphasize here. Since $\nu \in N_\beta$ must be nonzero to be a circuit, all circuits have *exactly* one negative entry. It is convenient for us to enumerate $\nu^+ := \{\alpha : \nu_\alpha > 0\}$, and to identify the unique index $\nu^- := \beta \in \mathcal{A}$ where $\nu_\beta < 0$. Separately, positive homogeneity of the support function tells us that the property of being a circuit is invariant under scaling by positive constants. A circuit is *normalized* if its unique negative term ν_β has $\nu_\beta = -1$, in which case we usually denote the circuit by “ λ .” Note that a circuit can be normalized by taking the ratio with its infinity norm $\nu/\|\nu\|_\infty$, because $\|\nu\|_\infty = |\nu_\beta|$.

In the special case when X is a convex cone, it is straightforward to determine which $\nu \in N_\beta$ are (\mathcal{A}, X) -circuits. In such a setting we have $\sigma_X(-\mathcal{A}\nu) < +\infty$ if and only if $\sigma_X(-\mathcal{A}\nu) = 0$, therefore $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ is trivially linear over the entirety of $V_\beta := \{\nu \in N_\beta : \sigma_X(-\mathcal{A}\nu) < +\infty\}$. We can derive an expression for V_β by reformulating the condition $\sigma_X(-\mathcal{A}\nu) = 0$ as $\nu \in (-\mathcal{A}^T X)^\circ = (\mathcal{A}^T X)^*$. The result of this process is that

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

where \mathcal{A}^+ denotes the pseudo-inverse of $\mathcal{A} : \mathbb{R}^{\mathcal{A}} \rightarrow \mathbb{R}^n$. Thus we see that when X is a cone, circuits $\nu \in N_\beta$ are precisely the edge generators of $(\ker \mathcal{A} + \mathcal{A}^+ X^*) \cap N_\beta$.

Example 3.3. Suppose $X = \mathbb{R}^n$. Then $X^* = \{0\}$, so $\ker \mathcal{A} + \mathcal{A}^+ X^* = \ker \mathcal{A}$ and $V_\beta = \ker \mathcal{A} \cap N_\beta$. It is easily shown that edge generators of $\ker \mathcal{A} \cap N_\beta$ are precisely those $\nu \in \ker \mathcal{A} \cap N_\beta \setminus \{0\}$ for which $\nu^+ = \{\alpha : \nu_\alpha > 0\}$ are affinely independent. Thus, Definition 3.1 recovers the matroid-theoretic notion of affine-linear simplicial circuits.

The following proposition shows that the affine-independence property is a necessary condition for *all* X -circuits, regardless of X . 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 constant (identically equal to $\sigma_X(z)$) on U , and so in order for ν^* to be an X -circuit, ν^* must be a vertex of the polytope U . The set U is in 1-to-1 correspondence with $W = \{w \in \mathbb{R}^{\mathcal{A} \setminus \beta} : \sum_{\alpha \in \mathcal{A} \setminus \beta} (\beta - \alpha)w_\alpha = z, 1^T w = -\nu_\beta^*\}$ by identifying $w = \nu_{\setminus \beta}$.

One may verify that affine-independence of a point set $Y \subset \mathbb{R}^\ell$ is translation and reflection invariant, as well as invariant under liftings of the form $Y' = Y \times \{1\}$. Basic polyhedral geometry tells us that all vertices w^* of W use an affinely independent set of columns $\{(\beta, 1) - (\alpha, 1) : w_\alpha^* > 0\}$ from the matrix in the linear system $[(\beta - \mathcal{A}) \times \{1\}]w = (z, -\nu_\beta^*)$. 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 $\alpha_1 = (0, 0)^T$, $\alpha_2 = (1, 0)^T$, and $\alpha_3 = (0, 1)^T$, and consider $X = \{x \in \mathbb{R}^2 : x \geq u\}$ for some fixed point $u \in \mathbb{R}^2$. The vector $\nu^* = (-2, 1, 1)^T$ has $(\nu^*)^- = \alpha_1 = (0, 0)^T$, and $(\nu^*)^+ = \{\alpha_2, \alpha_3\} = \{(1, 0)^T, (0, 1)^T\}$ is affinely independent. Considering $\nu^{(1)} = (-2, 2, 0)^T$ and $\nu^{(2)} = (-2, 0, 2)^T$, we have $\nu^* = \frac{1}{2}(\nu^{(1)} + \nu^{(2)}) \in \text{ri } L$ for $L := [\nu^{(1)}, \nu^{(2)}]$, where ri denotes the relative interior. 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 $(1, 1)^T$ is as well an optimal point of both the objective functions $x \mapsto (-2\mu_1, 0)^T x$ and $x \mapsto (0, -2\mu_2)^T x$ on X .

The fact that Example 3.5 took $X = u + \mathbb{R}_+^2$ for a symbolic vector $u \in \mathbb{R}^2$ illustrates a very important fact: the property of being an X -circuit is invariant under translation of X .

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 (\mathcal{A}, X) -circuit 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 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 [20, 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 a circuit in order for $(\nu^*, \sigma_X(-\mathcal{A}\nu^*))$ to be an edge generator of $\text{co}T$. The harder direction is showing that ν^* being a 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 $\Delta_\ell := \{z \in \mathbb{R}_+^\ell : 1^T z = 1\}$, 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 $F \trianglelefteq H$ containing 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 a 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 “ $\text{co}T$,” and hence only finitely many normalized circuits. The remainder of this section serves to prove this fact; here we use the concept of *normal fans* from polyhedral geometry. For each face F of a polyhedron P , there is an associated *outer normal cone*

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

The essential property of such a normal cone is that $w \mapsto \sigma_P(w)$ is finite and linear on $w \in \mathbf{N}_P(F)$, 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:

$$\mathcal{O}(P) = \{\mathbf{N}_P(F) : F \trianglelefteq P\}.$$

Normal fans are typically encountered in the study of polytopes (see e.g. [21, Chapter 7]) but their most important properties apply equally well to unbounded polyhedra.

Theorem 3.7. *If X is polyhedral and $\nu \in N_\beta$ is an (\mathcal{A}, X) -circuit, then $\text{co}\{\nu\}$ is a ray in $\mathcal{O}(P)$ for $P := -\mathcal{A}^T X + N_\beta^\circ$. Consequently, polyhedral X have finitely many normalized circuits.*

Proof. Let $P = -\mathcal{A}^T X + N_\beta^\circ$. Using the characterization in [20, 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 $F \trianglelefteq P$, 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 . It is known ([22, Corollary 1]) that the cones $K \in \mathcal{O}$ partition $(\text{rec } P)^\circ$, i.e., that

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

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} which contains ν in its relative interior. If K is of dimension greater than 1, then ν can be expressed as a convex combination of non-proportional $\nu^{(1)}, \nu^{(2)} \in K$ – and clearly $\nu \mapsto \sigma_X(-\mathcal{A}\nu) \equiv \sigma_P(\nu)$ 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 circuits associated with (\mathcal{A}, X) . \square

4. X -CIRCUITS IN AGE CONES

This section lays the foundation for using “ (\mathcal{A}, X) -circuits” to understand the extreme rays of the conditional SAGE cone $C_X(\mathcal{A})$. Our first result here is Theorem 4.2, which shows that if $\nu \neq \mathbf{0}$ is not an X -circuit but certifies nonnegativity of some f by way of (1), then f cannot be extremal in the corresponding X -AGE cone. This result is combined with Theorem 3.7 to obtain Theorem 4.4, which states that for polyhedral X , extremal AGE functions f admit *exactly* one ν certifying (1). Our focus then shifts to *normalized circuits*: λ satisfying Definition 3.1 with unique negative component $\lambda_\beta = -1$. Theorem 4.6 formalizes how to use normalized circuits to construct $C_X(\mathcal{A})$, and is crucial for subsequent development in Section 5. Corollaries 4.7 and 4.8 leverage Theorem 4.6 to produce representation results when X is a polyhedron. Theorem 4.6 and Proposition 4.10 are equally important for the development in Section 5.

The following lemma provides a construction to decompose an 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 $f(x) = \sum_{\alpha \in \mathcal{A}} c_\alpha \exp(\alpha^T x)$ be an X -AGE function with negative term $c_\beta < 0$. If ν satisfying (1) can be written as a convex combination $\nu = \sum_{i=1}^k \theta_i \nu^{(i)}$ of non-proportional $\nu^{(i)}$ 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

$$c_\alpha^{(i)} = \begin{cases} (c_\alpha/\nu_\alpha)\nu_\alpha^{(i)} & \text{if } \alpha \in \nu^+ \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } \alpha \in \mathcal{A} \setminus \beta, \quad (2)$$

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 (2) 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

$$\sigma_X(-\mathcal{A}\nu) = \sigma_X(-\mathcal{A}(\sum_{i=1}^k \theta_i \nu^{(i)})) = \sum_{i=1}^k \theta_i \sigma_X(-\mathcal{A}\nu^{(i)}). \quad (3)$$

Meanwhile, (2) 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

$$\sum_{i=1}^k \theta_i D(\nu_{\setminus\beta}^{(i)}, ec_{\setminus\beta}^{(i)}) = D(\nu_{\setminus\beta}, ec_{\setminus\beta}). \quad (4)$$

We combine (3) and (4) to obtain the desired result

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

Theorem 4.2. *Let $f(x) = \sum_{\alpha \in \mathcal{A}} c_\alpha \exp(\alpha^T x)$ be X -AGE with negative term $c_\beta < 0$. If $\nu \in \mathbb{R}^{\mathcal{A}}$ satisfies (1) but is not an X -circuit, then f is not extremal in $C_X(\mathcal{A}, \beta)$.*

Proof. If f is an AGE function with $c_\beta < 0$ and ν satisfies (1), then we must have $\nu \neq 0$ 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. \square

Our next lemma shows that extremal X -AGE functions in $C_X(\mathcal{A}, \beta)$ have a unique scale for ν satisfying (1). It is a prelude to the latter half of this section, which deals exclusively with normalized circuits.

Lemma 4.3. *If $f(x) = \sum_{\alpha \in \mathcal{A}} c_\alpha \exp(\alpha^T x)$ is X -AGE with $c_\beta < 0$ and there exist two distinct but proportional vectors $\nu^{(1)}, \nu^{(2)}$ satisfying (1), then f is not extremal in $C_X(\mathcal{A}, \beta)$.*

Proof. Define $F_c(\nu) = \sigma_X(-\mathcal{A}\nu) + D(\nu_{\setminus\beta}, ec_{\setminus\beta}) - c_\beta$. By assumption we have $F_c(\nu^{(i)}) \leq 0$, and since $c_\beta < 0$, neither $\nu^{(i)}$ can be zero. First consider the case that one of the $\nu^{(i)}$, say, $\nu^{(1)}$, gives a negative value for F_c , i.e., $F_c(\nu^{(1)}) < 0$. Then we set $\hat{c}_\beta = c_\beta + F_c(\nu^{(1)}) < c_\beta$. The

signomial f can be decomposed as $f = \hat{f} + p$ where $\hat{f}(x) = \sum_{\alpha \in \mathcal{A} \setminus \beta} c_\alpha \exp(\alpha^T x) + \hat{c}_\beta \exp(\beta^T x)$ and $p(x) = -F_c(\nu^{(1)}) \exp(\beta^T x)$ is a posynomial. Since \hat{f} is X -AGE with certificate $\nu^{(1)}$ by construction, the decomposition $f = \hat{f} + p$ certifies that f is nonextremal in $C_X(\mathcal{A})$.

It remains to address the possibility of $F_c(\nu^{(1)}) = F_c(\nu^{(2)}) = 0$. For fixed $\bar{\nu}$, the mapping $t \mapsto F_c(t\bar{\nu})$ for $t > 0$ is the sum of a linear, a strictly convex and a constant mapping, hence it is strictly convex. Since $\nu^{(1)}$ and $\nu^{(2)}$ are distinct nonzero proportional vectors, this rules out the case $F_c(\nu^{(1)}) = F_c(\nu^{(2)}) = 0$. \square

Theorem 4.4. *Let X be a polyhedron. Suppose $f(x) = \sum_{\alpha \in \mathcal{A}} c_\alpha \exp(\alpha^T x)$ is X -AGE with $c_\beta < 0$, and that more than one ν satisfies (1). Then f is not extremal in $C_X(\mathcal{A}, \beta)$.*

Proof. Let $\nu^{(1)}, \nu^{(2)}$ denote distinct vectors where $(\nu^{(1)}, c), (\nu^{(2)}, c)$ satisfy (1). If these vectors are proportional to one another, Lemma 4.3 implies the claim, therefore we may assume that $\nu^{(i)}$ are not proportional to one another. By convexity, we have that $\nu = \lambda \nu^{(1)} + (1 - \lambda) \nu^{(2)}$ has (ν, c) satisfy (1) for all $0 \leq \lambda \leq 1$. There are infinitely many such values of ν , but by Theorem 3.7 there are only finitely many X -circuits. Thus f admits an X -AGE certificate (ν, c) where ν is not an X -circuit. The result follows by Theorem 4.2. \square

In the remainder of this section we show the benefit of eliminating the degree of freedom associated with ν laying on a ray. To describe our results, we use

$$\Lambda_X(\mathcal{A}, \beta) = \{\lambda \in N_\beta : \lambda \text{ is an } X\text{-circuit (per Definition 3.1), } \lambda_\beta = -1\} \quad (5)$$

to denote normalized X -circuits associated with $\beta \in \mathcal{A}$. It also helps to aggregate

$$\Lambda_X(\mathcal{A}) := \bigcup_{\beta \in \mathcal{A}} \Lambda_X(\mathcal{A}, \beta). \quad (6)$$

Remark 4.5. One may verify that a vector $\lambda^* \in N_\beta$ belongs to $\Lambda_X(\mathcal{A}, \beta)$ if and only if (1) $\lambda_\beta^* = -1$, (2) $\sigma_X(-\mathcal{A}\lambda^*) < +\infty$ and (3) if $\lambda \mapsto \sigma_X(-\mathcal{A}\lambda)$ is linear on a line segment $[\lambda^{(1)}, \lambda^{(2)}]$ where $\lambda^{(i)} \in N_\beta$ are distinct with $\lambda_\beta^{(i)} = -1$, then $\lambda^* \notin \text{ri}[\lambda^{(1)}, \lambda^{(2)}]$.

We now define λ -witnessed AGE cones

$$C_X(\mathcal{A}, \lambda) = \left\{ c \in \mathbb{R}^{\mathcal{A}} : \beta := \lambda^-, \prod_{\alpha \in \lambda^+} \left[\frac{c_\alpha}{\lambda_\alpha} \right]^{\lambda_\alpha} \geq -c_\beta \exp(\sigma_X(-\mathcal{A}\lambda)), c_{\setminus \beta} \geq 0 \right\}. \quad (7)$$

The term ‘‘witnessed’’ in ‘‘ λ -witnessed AGE cone’’ is chosen because (as Theorem 4.6 will show) λ acts as a nonnegativity certificate for signomials with coefficient vectors $c \in C_X(\mathcal{A}, \lambda)$.

Theorem 4.6. $C_X(\mathcal{A}, \beta) = \text{conv} \bigcup_{\lambda \in \Lambda_X(\mathcal{A}, \beta)} C_X(\mathcal{A}, \lambda)$.

Proof. Theorem 4.2 already tells us that $C_X(\mathcal{A}, \beta)$ may be expressed as the convex hull of X -AGE functions $f(x) = \sum_{\alpha \in \mathcal{A}} c_\alpha \exp(\alpha^T x)$ where (c, ν) satisfies (1) for some X -circuit ν . Therefore it suffices to show that (i) for any such function, the normalized X -circuit $\lambda = \nu / (1^T \nu_{\setminus \beta})$ is such that (c, λ) satisfy the condition in (7), and (ii) if any (c, λ) satisfy (7), 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 (7).

Let $r(\nu) = \sigma_X(-\mathcal{A}\nu) + D(\nu_{\setminus\beta}, c_{\setminus\beta}) - 1^T \nu_{\setminus\beta}$. Apply positive homogeneity of the support function to see $\sigma_X(-\mathcal{A}\nu) = (1^T \nu_{\setminus\beta}) \sigma_X(\mathcal{A}\nu / (1^T \nu_{\setminus\beta}))$, and use $\nu = s\lambda$ to simplify $\sigma_X(-\mathcal{A}\nu / (1^T \nu_{\setminus\beta})) = \sigma_X(-\mathcal{A}\lambda)$. Abbreviate $d := \sigma_X(-\mathcal{A}\lambda)$, so as to express

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

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

$$-c_\beta \leq \prod_{\alpha \in \lambda^+} [\tilde{c}_\alpha / \lambda_\alpha]^{\lambda_\alpha}. \quad (8)$$

Since $[\tilde{c}_\alpha / \lambda_\alpha]^{\lambda_\alpha} = [c_\alpha / \lambda_\alpha]^{\lambda_\alpha} (\exp(-d))^{\lambda_\alpha}$ and $\prod_{\alpha \in \lambda^+} (\exp(-d))^{\lambda_\alpha} = \exp(-d)$, (8) can be recognized as the inequality occurring within (7), which completes the proof. \square

Theorem 4.6 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

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

the corresponding dual cone is given by

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

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 two corollaries concerning power-cone representability of $C_X(\mathcal{A})$ when X is a polyhedron. The first one generalizes the case $X = \mathbb{R}^n$ considered by Papp for polynomials [23].

Corollary 4.7. *If X is a polyhedron, then $C_X(\mathcal{A})$ is power-cone representable.*

Proof. 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 result follows as each of the finitely many sets $C_X(\mathcal{A}, \lambda)$ appearing in the above sum are (dual) power-cone representable. \square

The following corollary generalizes results by Averkov [24] and Wang and Magron [25] for ordinary SAGE polynomials, and recent results by Naumann and Theobald for several types of ordinary SAGE-like certificates [26].

Corollary 4.8. *Suppose X is a rational polyhedron and each $\alpha \in \mathcal{A}$ is a rational vector. Then Averkov’s “semidefinite extension degree [24]” of $C_X(\mathcal{A})$ is two.*

Proof. We observe that under the given rationality assumptions, each of the finitely many X -circuits will be rational. 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-cone representable when $\lambda_{\setminus\beta}$ is a rational vector in the $(m - 1)$ -dimensional probability simplex [27, Section 3.4]. The semidefinite extension degree of the second-order-cone is two, so the claim follows by Corollary 4.7. \square

To express $C_X(\mathcal{A}, \lambda)$ in terms of $\text{Pow}(\lambda)^*$, introduce a diagonal linear operator $S_\lambda : \mathbb{R}^{\mathcal{A}} \rightarrow \mathbb{R}^{\text{supp } \lambda}$ 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 $\beta := \lambda^-$. Also, let $\delta_\beta \in \mathbb{R}^{\mathcal{A}}$ denote the standard basis vector corresponding to $\beta \in \mathcal{A}$, i.e. $\delta_\beta^T w = w_\beta$ for $w \in \mathbb{R}^{\mathcal{A}}$.

Proposition 4.9. *For an X -circuit $\lambda \in \Lambda_X(\mathcal{A})$, the λ -witnessed AGE cone can be given by*

$$C_X(\mathcal{A}, \lambda) = \{c \in \mathbb{R}^{\mathcal{A}} : \beta := \lambda^-, c_{\setminus\beta} \geq 0, (S_\lambda c - r\delta_\beta) \in \text{Pow}(\lambda)^*, r \geq 0\}. \quad (9)$$

Proof. First, we note that some inequality constraints $c_{\setminus\beta} \geq 0$ 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

$$\prod_{\alpha \in \lambda^+} [c_\alpha / \lambda_\alpha]^{\lambda_\alpha} \geq |c_\beta \exp(\sigma_X(-\mathcal{A}\lambda)) - r|. \quad (10)$$

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 $r = c_\beta$ when $c_\beta \geq 0$. In the $c_\beta < 0$ case the constraint (10) 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 (10) is vacuous, since $\prod_{\alpha \in \lambda^+} [c_\alpha / \lambda_\alpha]^{\lambda_\alpha} \geq 0$ is implied by $c_{\setminus\beta} \geq 0$. As the constraint in the preceding display is similarly vacuous when $c_\beta > 0$, we see that it can be used in lieu of (10) without loss of generality. \square

We can appeal to Proposition 4.9 to find a representation for $C_X(\mathcal{A}, \lambda)^*$ which is analogous to Equation 7.

Proposition 4.10. *For an X -circuit $\lambda \in \Lambda_X(\mathcal{A})$, the dual λ -witnessed AGE cone is*

$$C_X(\mathcal{A}, \lambda)^* = \left\{ v \in \mathbb{R}_+^{\mathcal{A}} : \beta := \lambda^-, \exp(\sigma_X(-\mathcal{A}\lambda)) \prod_{\alpha \in \lambda^+} v_\alpha^{\lambda_\alpha} \geq v_\beta \right\}. \quad (11)$$

Proof. Let $\beta = \lambda^-$ as is usual. To $v \in \mathbb{R}^{\mathcal{A}}$ 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 assume for ease of exposition that $\mathcal{A} = \text{supp } \lambda$.³ When considering the given expression for $\text{Val}(v)$ as a primal problem, we compute a dual using (9) from Proposition 4.9. Under the

³If $\alpha \in \mathcal{A} \setminus \text{supp } \lambda$, the only constraints on c_α, v_α for $c \in C_X(\mathcal{A}, \lambda), v \in C_X(\mathcal{A}, \lambda)^*$ are $c_\alpha \geq 0, v_\alpha \geq 0$.

assumption $\mathcal{A} = \text{supp } \lambda$, the constraint $c_{\setminus\beta} \geq 0$ is implied by $(S_\lambda c - r\delta_\beta) \in \text{Pow}(\lambda)^*$. Therefore when forming a Lagrangian for $\text{Val}(v)$ using (9), the dual variable to “ $c_{\setminus\beta} \geq 0$ ” 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 X -CIRCUITS IN SAGE CONES

The previous section showed that an X -SAGE cone is generated by X -circuits. In the case $X = \mathbb{R}^n$ it is known that every simplicial circuit “ λ ” generates a weighted AGE cone $C_{\mathbb{R}^n}(\mathcal{A}, \lambda)$ containing an extreme ray of $C_{\mathbb{R}^n}(\mathcal{A}, \beta)$. However, Katthän, Naumann, and Theobald showed that many extreme rays of AGE cones are *not* extreme when considered in the sum $C_{\mathbb{R}^n}(\mathcal{A}) = \sum_{\beta \in \mathcal{A}} C_{\mathbb{R}^n}(\mathcal{A}, \beta)$ [7]. Specifically, a normalized \mathbb{R}^n -circuit $\lambda \in \Lambda_{\mathbb{R}^n}(\mathcal{A})$ is only needed in $C_{\mathbb{R}^n}(\mathcal{A})$ if $(\text{ri conv}\{\alpha\}_{\alpha \in \lambda^+}) \cap \mathcal{A} = \{\lambda^-\}$ [7, Proposition 4.4]. Circuits satisfying this property were said to be *reduced*. The idea of reduced circuits was similarly considered by Forsgård and de Wolff, who defined the *Reznick cone* associated with \mathcal{A} as the conic hull $R_{\mathbb{R}^n}(\mathcal{A}) := \text{co } \Lambda_{\mathbb{R}^n}(\mathcal{A})$ [8]. Theorem 3.2 of [8] says a vector λ is an edge generator of $R_{\mathbb{R}^n}(\mathcal{A})$ if and only if $(\text{ri conv}\{\alpha\}_{\alpha \in \lambda^+}) \cap \mathcal{A} = \{\lambda^-\}$. Combining this theorem with Katthän, Naumann, and Theobald’s Proposition 4.4, one can say that the extreme rays of the SONC or \mathbb{R}^n -SAGE cone are supported on circuits which are edge generators of $R_{\mathbb{R}^n}(\mathcal{A})$.⁴

We shall prove an analogous result in Theorem 5.5. Per Remark 3.2, our use of the term “circuit” should be understood to mean (\mathcal{A}, X) -circuit, where \mathcal{A} and X are clear from context. Throughout this section we make appeals to convex geometry which necessitate moving back and forth between normalized and unnormalized circuits. The set of normalized circuits $\Lambda_X(\mathcal{A})$ is bounded and nonconvex. When we have need to enumerate unnormalized circuits, we simply refer to “ ν with $\nu/\|\nu\|_\infty \in \Lambda_X(\mathcal{A})$.”

⁴Forsgård and de Wolff provide this statement without proof following their Theorem 3.2.

Definition 5.1. The *functional form* of a circuit ν with $\nu/\|\nu\|_\infty \in \Lambda_X(\mathcal{A})$ is $\phi_\nu : \mathbb{R}^A \rightarrow \mathbb{R}$ 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 $\phi_\nu = (\nu, \sigma_X(-\mathcal{A}\nu)) \in \mathbb{R}^A \times \mathbb{R}$ to denote the functional form of a given circuit. When representing the functional form of a circuit by a vector in $\mathbb{R}^A \times \mathbb{R}$, 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 graph* of (\mathcal{A}, X) is $G_X(\mathcal{A}) = \text{co}(\{\phi_\lambda : \lambda \in \Lambda_X(\mathcal{A})\} \cup \{(\mathbf{0}, 1)\})$, where $(\mathbf{0}, 1) \in \mathbb{R}^A \times \mathbb{R}$.

An equivalent construction is that $G_X(\mathcal{A})$ is the conic hull of the epigraph of $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$, as ν varies over the nonconvex set of unnormalized (\mathcal{A}, X) -circuits. We use the construction in Definition 5.2 for a few reasons. First, it has a more geometric flavor. In particular it is natural to think of the “essence” $G_X(\mathcal{A})$ as coming from the smaller cone $\text{co}\{\phi_\lambda : \lambda \in \Lambda_X(\mathcal{A})\}$. Second, highlighting the contribution of the vector $(\mathbf{0}, 1)$ makes it easier to reason about certain aspects of the *dual* circuit graph $G_X(\mathcal{A})^* \subset \mathbb{R}^A \times \mathbb{R}_+$. The dual circuit graph plays a central role in this section, particularly via the following theorem.

Theorem 5.3. $C_X(\mathcal{A})^* = \text{cl}\{\text{exp } y : (y, 1) \in G_X(\mathcal{A})^*\}$.

The circuit graph is very similar to the Reznick cone of Forsgård and de Wolff. The simplest distinction is that $G_X(\mathcal{A})$ is a subset of $\mathbb{R}^A \times \mathbb{R}$, rather than a subset of \mathbb{R}^A . The way in which the extension from \mathbb{R}^A to $\mathbb{R}^A \times \mathbb{R}$ affects $G_X(\mathcal{A})$ depends jointly on the structure of \mathcal{A} and X . When X is a cone (such as $X = \mathbb{R}^n$), any X -circuit ν has $\sigma_X(-\mathcal{A}\nu) = 0$, and so $G_X(\mathcal{A})$ is by all accounts equivalent to the projection $\text{co } \Lambda_X(\mathcal{A}) = \{\nu : (\nu, 0) \in G_X(\mathcal{A})\}$.

Definition 5.4. The *reduced circuits* of (\mathcal{A}, X) 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 $G_X(\mathcal{A})$.

We are mostly interested in reduced circuits which are normalized, and use $\Lambda_X^*(\mathcal{A})$ to denote the set of all such vectors. We shall prove in this section that Definition 5.4’s notion of a “reduced circuit” provides a construction of an X -SAGE cone $C_X(\mathcal{A})$ from fewer λ -witnessed AGE cones than are used in Theorem 4.2. In fact, the definition of a “circuit graph” was specifically formulated so that $\Lambda_X^*(\mathcal{A})$ provides the *unique minimal construction* of $C_X(\mathcal{A})$ from λ -witnessed AGE cones in the case when X is polyhedral. It might be instructive to consider the various statements proven in this section if $G_X(\mathcal{A})$ was defined as the smaller set $\text{co}\{\phi_\lambda : \lambda \in \Lambda_X(\mathcal{A})\}$, in which case the analogous definition of “reduced circuit” would be slightly more permissive.

With Definitions 5.1, 5.2, and 5.4 established, we can state the main results of this section.

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

$$C_X(\mathcal{A}) = \text{cl}\left(\text{conv}\bigcup\left\{C_X(\mathcal{A}, \lambda) : \lambda \in \Lambda_X^*(\mathcal{A})\right\}\right). \quad (12)$$

We point out how Theorem 5.5 involves a closure around the union over λ -witnessed AGE cones, while Theorem 4.6 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 4.6 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*

$$C_X(\mathcal{A}) = \sum_{\lambda \in \Lambda_X^*(\mathcal{A})} C_X(\mathcal{A}, \lambda). \quad (13)$$

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 actually 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 remainder of this section is organized as follows. Section 5.1 proves Theorem 5.3, which is instrumental in later subsections. In Section 5.2 we introduce and prove a certain representation result for the circuit graph. Given the groundwork laid in these first two subsections, Section 5.3 proves Theorem 5.5 in very short order. Section 5.4 proves Theorem 5.6 by refining the arguments from Section 5.3.

5.1. Proof of Theorem 5.3. Our proof of Theorem 5.3 uses the following simple lemma.

Lemma 5.7. *If $S \subset T$ are convex sets with $S \cap \text{ri} T \neq \emptyset$, then $S = \text{cl}(S \cap \text{ri} T)$.*

Proof. Rockafellar's [20, 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 $S = \text{cl} \text{ri} S$. \square

Proof of Theorem 5.3. Use Rockafellar's [20, Corollary 16.5.2] to invoke Theorem 4.6 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 4.10 implies

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

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 5.7 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 (14) 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 $(\mathbf{0}, 1)$ simply requires $(y, t) \in G_X(\mathcal{A})^*$ to have $t \geq 0$.

5.2. Topological properties of the circuit graph.

Theorem 5.8. $G_X(\mathcal{A}) = \text{co}(\{\phi_\lambda : \lambda \in \Lambda_X^*(\mathcal{A})\} \cup \{(\mathbf{0}, 1)\})$.

The proof of this theorem essentially reduces to showing that $G_X(\mathcal{A})$ is pointed and closed. The pointedness of the circuit graph 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 5.9. *The closure of the circuit graph 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 5.3 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 5.10. *The circuit graph 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 5.9), 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 5.8. Lemmas 5.9 and 5.10 show $G_X(\mathcal{A})$ is closed and pointed. By [20, 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

5.3. 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 5.8, we can express

$$(y, 1) \in G_X(\mathcal{A})^* \Leftrightarrow (y, 1)^T(\lambda, \sigma_X(-\mathcal{A}\lambda)) \geq 0 \quad \forall \lambda \in \Lambda_X^*(\mathcal{A}). \quad (15)$$

We obtain the following refinement of Equation 14, by combining (15) with Theorem 5.3:

$$C_X(\mathcal{A})^* = \left\{ v \in \mathbb{R}_+^{\mathcal{A}} : \forall \lambda \in \Lambda_X^*(\mathcal{A}), \beta := \lambda^-, \exp(\sigma_X(-\mathcal{A}\lambda)) \prod_{\alpha \in \lambda^+} v_\alpha^{\lambda_\alpha} \geq v_\beta \right\}. \quad (16)$$

Of course, Equation 16 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, [20, Corollary 16.5.2]) to obtain the claim of the theorem. \square

5.4. 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 which 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 5.11. *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, Theorem 3.7 tells us $\Lambda_X(\mathcal{A})$ is finite, so $\Lambda_X(\mathcal{A})$ is closed and $\Lambda_X^*(\mathcal{A})$ is finite. From closedness of $\Lambda_X(\mathcal{A})$ we have $G_X(\mathcal{A}) = \text{co } T_1$, and in particular every $\phi_\lambda \in$

$T_1 \setminus \{(\mathbf{0}, 1)\}$ is known to generate 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 $K = \mathbb{R}^A \times \mathbb{R}_+$. Furthermore, since X is presumed nonempty, Theorem 5.3 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

Our next lemma shows how to take a condition stated in terms of Lemma 5.11, and deduce a statement about $C_X(\mathcal{A})^*$. The lemma's proof requires only that X be nonempty and convex.

Lemma 5.12. *If $\tilde{y} \in \mathbb{R}^A$ 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 $z \in \mathbb{R}^A$ 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 5.3 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

$$\exp(\phi_\lambda(\tilde{y})) < 1 \leq \exp(\phi_\lambda(y)) \quad \text{for all } \exp y \in C_X(\mathcal{A})^*. \quad (17)$$

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 (17) by $\exp(-u - \tilde{y}_\beta + y_\beta)$ to obtain

$$0 \leq \exp\left(\lambda_{\setminus\beta}^T(y_{\setminus\beta} - \tilde{y}_{\setminus\beta})\right) - \exp(-u - \tilde{y}_\beta + y_\beta). \quad (18)$$

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 (18) gives

$$0 \leq \left(\lambda_{\setminus\beta} \circ \exp(-\tilde{y}_{\setminus\beta})\right)^T \exp(y_{\setminus\beta}) - (\exp(-u - \tilde{y}_\beta)) \exp(y_\beta). \quad (19)$$

Inequality (19) is essentially what we need to prove the lemma. Defining $z \in \mathbb{R}^A$ 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 $\lambda_{\setminus\beta}^T \mathbf{1} = 1$ 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 5.3, we have the dual description $C_X(\mathcal{A})^* = \text{cl}\{\text{exp } y : (y, 1) \in G_X(\mathcal{A})^*\}$. Applying Theorem 5.8 then gives

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

We rewrite the condition on $\phi_\lambda(y)$ as a condition on $v = \text{exp } y$ using the power-cone formulation in Proposition 4.10. 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 $\text{co}\{\text{exp}(\mathcal{A}^T x) : x \in X\}$, 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 [20, 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. $C^* \supsetneq C_X(\mathcal{A})^*$. 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 5.7 we have $C^* = \text{cl}(C^* \cap \mathbb{R}_{++}^A)$. Work with C^* over the positive orthant using Proposition 4.10 to express it as $C^* = \text{cl}\{\text{exp } y : y \in Y\}$ for $Y := \{y : \phi_\lambda(y) \geq 0 \ \forall \lambda \in \Lambda\}$. By Lemma 5.11 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 5.12 to this pair (ϕ_λ, y) to see that $\text{exp } \tilde{y}$ can be separated from the closed convex set $C_X(\mathcal{A})^*$. We have therefore found a point \tilde{y} where $\text{exp } \tilde{y} \in C^*$ and yet $\text{exp } \tilde{y}$ can be separated from $C_X(\mathcal{A})^*$, so we conclude $C^* \supsetneq C_X(\mathcal{A})^*$. \square

6. REDUCED CIRCUITS AND EXTREME RAYS FROM A PRIMAL PERSPECTIVE

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. However, the duality argument is not without its drawbacks. For one thing, Theorem 5.5 currently requires a closure around the expression for $C_X(\mathcal{A})$ – this creates an obstruction to using extremality of $\phi_\lambda \in G_X(\mathcal{A})$ as a necessary condition for extremality of λ -witnessed AGE functions. As a second side effect of the duality argument, it is not clear how one can decompose a known nonextremal X -AGE function into a sum of extremal X -AGE functions. This section means to ameliorate such limitations by approaching reduced circuits and extreme rays in a constructive manner; the results of this section are Theorem 6.1 and Proposition 6.3.

Theorem 6.1. Fix $\lambda \in \Lambda_X(\mathcal{A}, \beta)$. Assume there exists $\beta' \in \mathcal{A} \setminus \text{supp } \lambda$ and $\lambda' \in \Lambda_X(\mathcal{A}, \beta')$ where $(\lambda')^+ \subset \lambda^+$ and $\mathcal{A}\lambda' = \mathcal{A}\lambda$. Then, every X -AGE function f with nonnegativity witness λ can be decomposed into distinct X -AGE functions which are not multiples of each other.

Theorem 6.1 can be proven in two ways. Given the machinery developed in the article thus far, the shorter proof is to show that if (λ, λ') satisfy the stated condition, then ϕ_λ cannot be an edge generator for $G_X(\mathcal{A})$. Thus at least in the case of polyhedral X , Theorem 5.6 would imply that a λ -witnessed AGE function f could not be extremal in $C_X(\mathcal{A})$. The longer proof is to explicitly provide the decomposition of f into nonproportional AGE functions; this strategy has the benefit of being applicable to *any* convex set X , and highlights how cancellation of coefficients across X -AGE cones is a fundamental issue when studying the boundary of $C_X(\mathcal{A})$. Before diving into the decomposition proof in earnest, we establish the following proposition.

Proposition 6.2. Fix $\lambda \in \Lambda_X(\mathcal{A}, \beta)$. Assume there exists $\beta' \in \mathcal{A} \setminus \text{supp } \lambda$ and $\lambda' \in \Lambda_X(\mathcal{A}, \beta')$ where $(\lambda')^+ \subset \lambda^+$ and $\mathcal{A}\lambda' = \mathcal{A}\lambda$. Then there exist vectors $\hat{\lambda} \in N_\beta, \hat{\lambda}' \in N_{\beta'}$ for which

$$\lambda \in \text{ri}[\hat{\lambda}, \hat{\lambda}'] \quad \text{and} \quad \mathcal{A}\hat{\lambda} = \mathcal{A}\hat{\lambda}' = \mathcal{A}\lambda.$$

Proof. Let such λ, λ' be given. The proof proceeds by identifying certain vectors $\lambda^{(1)}, \lambda^{(2)}$ which end up being scalar multiples of the desired $\hat{\lambda}, \hat{\lambda}'$. The $\lambda^{(j)}$ are set to satisfy

$$\lambda^{(1)} = \lambda - \lambda_{\beta'}^{(1)} \lambda' \quad \text{and} \quad \lambda^{(2)} = \lambda' - \lambda_{\beta}^{(2)} \lambda;$$

the precise values for $\lambda^{(j)}$ are determined by taking $\lambda_{\beta'}^{(1)}$ and $\lambda_{\beta}^{(2)}$ as large as possible so that $\lambda^{(1)} \in N_\beta$ and $\lambda^{(2)} \in N_{\beta'}$. The crucial properties of this construction follow from $(\lambda')^+ \subset \lambda^+$: we have $0 < \lambda_{\beta'}^{(1)} < 1$ and $0 \leq \lambda_{\beta}^{(2)} < 1$. Next, define $\tau_1 = 1/(1 - \lambda_{\beta'}^{(1)} \lambda_{\beta}^{(2)})$ and $\tau_2 = \lambda_{\beta'}^{(1)}/(1 - \lambda_{\beta'}^{(1)} \lambda_{\beta}^{(2)})$. These values for τ_j are chosen to utilize the identity $\lambda^{(1)} + \lambda_{\beta'}^{(1)} \lambda^{(2)} = (1 - \lambda_{\beta'}^{(1)} \lambda_{\beta}^{(2)})\lambda$ between $\lambda^{(1)}, \lambda^{(2)}$, and λ . Specifically, direct calculations show

$$\lambda = \tau_1 \lambda^{(1)} + \tau_2 \lambda^{(2)} \quad \text{and} \quad (\tau_1(1 - \lambda_{\beta'}^{(1)})) + (\tau_2(1 - \lambda_{\beta}^{(2)})) = 1.$$

Taking the above as given, we see that λ is given by a convex combination of $\hat{\lambda} = \lambda^{(1)}/(1 - \lambda_{\beta'}^{(1)})$ and $\hat{\lambda}' = \lambda^{(2)}/(1 - \lambda_{\beta}^{(2)})$. The proposition follows upon verifying $\mathcal{A}\hat{\lambda} = \mathcal{A}\hat{\lambda}' = \mathcal{A}\lambda$. \square

Proposition 6.2 is stated in a way that if (λ, λ') satisfy its hypothesis, then nonreducedness of λ should be clear. Our proof of Theorem 6.1 uses the construction provided by Proposition 6.2. However, rather than using $(\hat{\lambda}, \hat{\lambda}')$ directly, the proof of Theorem 6.1 relies on the *unscaled* vectors $(\lambda^{(1)}, \lambda^{(2)})$. It is therefore important to confirm the properties of $(\lambda^{(1)}, \lambda^{(2)})$ from the proof of Proposition 6.2 before attempting to verify the following proof in detail.

Proof of Theorem 6.1. Let $f(x) = \sum_{\alpha \in \lambda^+} c_\alpha \exp(\alpha^T x) + c_\beta \exp(\beta^T x)$. Without loss of generality, we can assume that for $\lambda \in \Lambda_X(\mathcal{A}, \beta)$, the coefficients satisfy

$$c_\beta = - \prod_{\alpha \in \lambda^+} [c_\alpha / \lambda_\alpha]^{\lambda_\alpha} \exp(-\sigma_X(-\mathcal{A}\lambda)). \quad (20)$$

If (20) were not satisfied with equality, we could decompose f into a sum of an X -AGE function satisfying (20) with equality, and a monomial with positive coefficient. Throughout the proof we abbreviate $z = -\mathcal{A}\lambda$.

Now consider the other circuit $\lambda' \in \Lambda_X(\mathcal{A}, \beta')$ with $(\lambda')^+ \subset \lambda^+$ and $-\mathcal{A}\lambda' = z$. Of course – since both λ and λ' are normalized – we have $\sum_{\alpha \in \lambda^+} \lambda_\alpha = \sum_{\alpha \in \lambda^+} \lambda'_\alpha = 1$. From λ' we construct the single, distinguished scalar

$$c'_{\beta'} := -\prod_{\alpha \in \lambda^+} [c_\alpha / \lambda_\alpha]^{\lambda'_\alpha} \exp(-\sigma_X(z)). \quad (21)$$

We proceed by finding coefficient vectors $c^{(1)}, c^{(2)}$ for AGE signomials f_1, f_2 where $f \in \text{ri co}\{f_1, f_2\}$ and in particular $c_{\beta'}^{(2)} = c'_{\beta'} < 0$. Such signomials will be non-proportional, since $c_{\beta'} = 0$.

The values for $c^{(1)}, c^{(2)}$ are specified in terms of the vectors $\lambda^{(1)}, \lambda^{(2)}$ from the proof of Proposition 6.2. For $\alpha \in \lambda^+$ we take $c_\alpha^{(j)} = c_\alpha(\lambda_\alpha^{(j)} / \lambda_\alpha)$, and for indices β, β' :

$$\begin{aligned} c_\beta^{(1)} &= c_\beta, & c_{\beta'}^{(1)} &= -\lambda_{\beta'}^{(1)} c'_{\beta'} \quad (> 0, \text{ since } c'_{\beta'} < 0 \text{ and } \lambda_{\beta'}^{(1)} > 0), \\ c_\beta^{(2)} &= -\lambda_\beta^{(2)} c_\beta \quad (\geq 0, \text{ since } c_\beta < 0 \text{ and } \lambda_\beta^{(2)} \geq 0), & c_{\beta'}^{(2)} &= c'_{\beta'}. \end{aligned}$$

Any remaining components $c_\alpha^{(j)}$ for $\alpha \notin (\lambda^+) \cup \{\beta, \beta'\}$ are set to $c_\alpha^{(j)} = 0$. We can prove nonnegativity of f_1 by demonstrating

$$\prod_{\alpha \in (\lambda^{(1)})^+} \left(\frac{c_\alpha^{(1)}}{\lambda_\alpha^{(1)}} \right)^{\lambda_\alpha^{(1)}} \exp(-\sigma_X(-\mathcal{A}\lambda^{(1)})) = -c_{\beta'}^{(1)}. \quad (22)$$

Verifying (22) involves some careful rewriting of the product indexed by $\alpha \in (\lambda^{(1)})^+$. Over the course of this rewriting, we split the product over $\alpha \in (\lambda^{(1)})^+$ into a product over $\alpha \in \lambda^+ \cup \{\beta'\}$ and invoke Equation 21 to write $-c'_{\beta'}$ as a similar product over $\alpha \in \lambda^+$. From there, we group the two products into a single product over $\alpha \in \lambda^+$, and simplify exponents by the relation $\lambda^{(1)} = \lambda - \lambda_{\beta'}^{(1)} \lambda'$. The following display carries out this procedure, where the final step uses $\sigma_X(-\mathcal{A}\lambda^{(1)}) = (1 - \lambda_{\beta'}^{(1)})\sigma_X(z)$ and Equation 20:

$$\begin{aligned} & \prod_{\alpha \in (\lambda^{(1)})^+} \left(\frac{c_\alpha^{(1)}}{\lambda_\alpha^{(1)}} \right)^{\lambda_\alpha^{(1)}} \exp(-\sigma_X(-\mathcal{A}\lambda^{(1)})) = \left[\prod_{\alpha \in \lambda^+} \left(\frac{c_\alpha}{\lambda_\alpha} \right)^{\lambda_\alpha^{(1)}} \right] \cdot \left[(-c'_{\beta'})^{\lambda_{\beta'}^{(1)}} \right] \exp(-\sigma_X(-\mathcal{A}\lambda^{(1)})) \\ &= \left[\prod_{\alpha \in \lambda^+} \left(\frac{c_\alpha}{\lambda_\alpha} \right)^{\lambda_\alpha^{(1)}} \right] \cdot \left[\left(\prod_{\alpha \in \lambda^+} \left(\frac{c_\alpha}{\lambda_\alpha} \right)^{\lambda'_\alpha} \right)^{\lambda_{\beta'}^{(1)}} (\exp(-\sigma_X(z)))^{\lambda_{\beta'}^{(1)}} \right] \cdot \exp(-\sigma_X(-\mathcal{A}\lambda^{(1)})) \\ &= \left(\prod_{\alpha \in \lambda^+} \left(\frac{c_\alpha}{\lambda_\alpha} \right)^{\lambda_\alpha^{(1)} + \lambda_{\beta'}^{(1)} \lambda'_\alpha} \right) \cdot \exp(-\sigma_X(z)) = -c_\beta = -c_{\beta'}^{(1)}. \end{aligned}$$

An entirely analogous calculation likewise shows f_2 is nonnegative; the only distinctions are that one should use $\lambda^{(2)} = \lambda' - \lambda_\beta^{(2)} \lambda$ instead of $\lambda^{(1)} = \lambda - \lambda_{\beta'}^{(1)} \lambda'$, and that one should verify the possibility of $\lambda_\beta^{(2)} = 0$ does not affect validity of the calculations.

The last step of the proof is to show that $f \in \text{rico}\{f_1, f_2\}$. We particularly show $f = \tau_1 f_1 + \tau_2 f_2$ for $\tau_1 = 1 / (1 - \lambda_{\beta'}^{(1)} \lambda_{\beta}^{(2)})$ and $\tau_2 = \lambda_{\beta'}^{(1)} / (1 - \lambda_{\beta'}^{(1)} \lambda_{\beta}^{(2)})$. These values of τ_1, τ_2 are from the proof of Proposition 6.2, and ensure $\tau_1 \lambda^{(1)} + \tau_2 \lambda^{(2)} = \lambda$. We claim that for all $\alpha \in \lambda^+ \cup \{\beta, \beta'\}$, we have $\tau_1 c_{\alpha}^{(1)} + \tau_2 c_{\alpha}^{(2)} = c_{\alpha}$. Working through the calculations: for $\alpha \in \lambda^+$

$$\tau_1 c_{\alpha}^{(1)} + \tau_2 c_{\alpha}^{(2)} = \frac{c_{\alpha}}{\lambda_{\alpha}} (\tau_1 \lambda_{\alpha}^{(1)} + \tau_2 \lambda_{\alpha}^{(2)}) = c_{\alpha},$$

and for index β

$$\tau_1 c_{\beta}^{(1)} + \tau_2 c_{\beta}^{(2)} = \frac{1}{1 - \lambda_{\beta'}^{(1)} \lambda_{\beta}^{(2)}} (1 - \lambda_{\beta'}^{(1)} \lambda_{\beta}^{(2)}) c_{\beta} = c_{\beta} \quad (\text{recall } 0 \leq \lambda_{\beta'}^{(1)}, \lambda_{\beta}^{(2)} < 1),$$

and finally index β' has $\tau_1 c_{\beta'}^{(1)} + \tau_2 c_{\beta'}^{(2)} = 0$. □

The last result of this article is to examine the 1-dimensional case $X = [0, \infty)$ in detail.

Proposition 6.3. *Extreme rays of $C_{[0, \infty)}(\mathcal{A})$ are exactly the following:*

- (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. Since SAGE cones are invariant under the application of a translation to all vectors in \mathcal{A} , we can assume $\alpha_1 = 0$ without loss of generality. 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 circuit. 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 AGE functions with two or three terms.

First we address the monomials. Given $f(x) = \exp(\alpha_i)$ 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 $f(x) = \exp(\alpha_1 x) = 1$. Since $X = [0, \infty)$, the leading term of any AGE function g must have a positive coefficient. Moreover, if g is not proportional to f , the leading term of g must be nonconstant, and so $\lim_{x \rightarrow \infty} g(x) = +\infty$. Since f is constant on $[0, \infty)$, we conclude f cannot be written as a convex combination of elements in $C_{[0, \infty)}(\mathcal{A})$ which are not proportional to itself, and so f is extremal in $C_{[0, \infty)}(\mathcal{A})$.

Now we consider the 2-term case. It is clear that $f(x) = c_j \exp(\alpha_j x) - c_i \exp(\alpha_i x)$ with $j > i$ is nonnegative on $[0, \infty)$ if and only if $c_j \geq c_i \geq 0$, and furthermore that such signomials are nonextremal unless $c_j = c_i$. Thus, when studying extremal 2-term signomials in $C_{[0, \infty)}(\mathcal{A})$, it suffices to consider $c_j = c_i = 1$ and $j > i$. We now show any such signomials with $j > i + 1$ or $i > 1$ are *nonextremal*. The case $j > i + 1$ is simple, as we have the decomposition

$f = f_1 + f_2$ for $f_1(x) = \exp(\alpha_j x) - \exp(\alpha_{j-1} x)$ and $f_2(x) = \exp(\alpha_{j-1} x) - \exp(\alpha_i x)$. Next, consider $j = i + 1$ and $i > 1$. Here we decompose $f = f_1 + f_2$ for parameterized functions $f_1(x) = b \exp(\alpha_i x) - b \exp(\alpha_{i-1} x)$ and $f_2(x) = \exp(\alpha_{i+1} x) - (1 + b) \exp(\alpha_i x) + b \exp(\alpha_{i-1} x)$ with $b \geq 0$. Clearly, f_1 is AGE; we claim that when b belongs to the nonempty open interval $(0, [\alpha_{i+1}/\alpha_i] - 1)$, the summand f_2 is also AGE. To see why, use absolute convergence of the Taylor series for $\exp(x)$ to write f_2 as

$$f_2(x) = \sum_{\ell=0}^{\infty} \frac{x^\ell}{\ell!} t_\ell \quad \text{for} \quad t_\ell := b(\alpha_{i-1})^\ell - (1 + b)\alpha_i^\ell + (\alpha_{i+1})^\ell.$$

Since $X = [0, \infty)$, nonnegativity of all t_ℓ suffices for nonnegativity of f_2 on X , and since $t_0 = 0$ we only need to consider $\ell \geq 1$. For this purpose use $\alpha_1 = 0$ and $i > 0$ to obtain scaled coefficients $s_\ell = t_\ell/(\alpha_i)^\ell$ with the same signs as t_ℓ . Then, bounding $s_\ell \geq -(1 + b) + (\alpha_{i+1}/\alpha_i)^\ell$, we use $\alpha_{i+1}/\alpha_i > 1$ and subsequently find $s_\ell \geq -(1 + b) + [\alpha_{i+1}/\alpha_i]$ for all $\ell \geq 1$. This shows that when $b < [\alpha_{i+1}/\alpha_i] - 1$, the coefficients s_ℓ and t_ℓ are nonnegative for all ℓ . By taking $b \in (0, [\alpha_{i+1}/\alpha_i] - 1)$ we thus decompose f as a sum of two non-proportional AGE functions f_1, f_2 . In closing, note $f(x) = \exp(\alpha_2 x) - \exp(\alpha_1 x)$ 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 proven cases (1) and (2) of the proposition. Now we consider 3-term AGE functions $f(x) = c_k \exp(\alpha_k x) + c_j \exp(\alpha_j x) + c_i \exp(\alpha_i x)$ for $k > j > i$. Start by considering the case when the trailing term $c_i < 0$. Given such f , construct $f_1(x) = (c_j + c_k) \exp(\alpha_j x) + c_i \exp(\alpha_i x)$ and $f_2(x) = c_k \exp(\alpha_k x) - c_k \exp(\alpha_j x)$. The summand f_2 is AGE since $c_k \geq 0$, and using $f(0) \geq 0 \Rightarrow |c_i| \leq c_j + c_k$, we see f_1 is also AGE. Therefore $c_i < 0$ implies f is nonextremal in $\mathcal{C}_{[0, \infty)}(\mathcal{A})$, and so all extremal 3-term AGE functions have middle term $c_j < 0$.

Using this knowledge of coefficient sign patterns for 3-term extremal AGE functions, we next prove that 3-term extremal AGE functions are nonnegative on the entirety of \mathbb{R} . To do this, let $\lambda \in \Lambda_X^*(\mathcal{A})$ with $\text{supp } \lambda = \{k, j, i\}$, $k > j > i$, and $\lambda^- = j$ be given; the claim will follow if we can show $\lambda \in \ker \mathcal{A}$. From λ we define $s = \min\{\lambda_k, \lambda_i\} > 0$, and construct the vector $\nu \in \mathbb{R}^m$ by $\nu_k = s(\alpha_j - \alpha_i)/(\alpha_k - \alpha_i)$, $\nu_j = -s$, $\nu_i = s(\alpha_k - \alpha_j)/(\alpha_k - \alpha_i)$, and $\nu_\ell = 0$ for $\ell \notin \{k, j, i\}$. This new vector ν satisfies $\mathcal{A}\nu = 0$, and since $X = [0, \infty)$, the original vector λ satisfies $\mathcal{A}\lambda \geq 0$. Given (λ, ν) , we construct another nonzero vector $\hat{\lambda} := \lambda - \nu/2$ so that $\hat{\lambda} \in N_j$ and $\mathcal{A}\hat{\lambda} \geq 0$. We therefore obtain a decomposition

$$\lambda = \frac{1}{2} (2\hat{\lambda}) + \frac{1}{2} \nu$$

showing λ belongs to the relative interior of a line segment $[2\hat{\lambda}, \nu] \subset N_j$ over which $\bar{\nu} \mapsto \sigma_X(-\mathcal{A}\bar{\nu})$ is linear. Therefore if $\lambda \in \Lambda_X^*(\mathcal{A})$, then λ is proportional to ν , and $\lambda \in \ker \mathcal{A}$.

At this point in the proof we can take advantage of an earlier result in the univariate case with $X = \mathbb{R}$: [7, Proposition 4.4]. Specifically, all $\lambda \in \Lambda_X^*(\mathcal{A}) \cap \ker \mathcal{A}$ satisfy $\lambda^+ = \{i + 1, i - 1\}$ and $\lambda^- = i$, and from these conditions one may solve for $\lambda_{i+1} = (\alpha_i - \alpha_{i-1})/(\alpha_{i+1} - \alpha_{i-1})$ and

$\lambda_{i-1} = (\alpha_{i+1} - \alpha_i)/(\alpha_{i+1} - \alpha_{i-1})$. Furthermore, it is known that every function of the form

$$f(x) = c_{i+1} \exp(\alpha_{i+1}x) - \left(\left[\frac{c_{i+1}}{\lambda_{i+1}} \right]^{\lambda_{i+1}} \left[\frac{c_{i-1}}{\lambda_{i-1}} \right]^{\lambda_{i-1}} \right) \exp(\alpha_i x) + c_{i-1} \exp(\alpha_{i-1}x) \quad (23)$$

cannot be expressed as a nontrivial convex combination of other functions of that same form (i.e. \mathbb{R} -AGE functions supported on \mathcal{A}) [7, Proposition 4.4].

We have arrived at the final phase of proving part (3) of this proposition. By appeal to the AM/GM inequality⁵, one finds the unique minimizer x^* for functions (23) is that satisfying

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

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 (23) 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 [7, Proposition 4.4], therefore any f given by (23) with $V_i(\lambda, c) > 1$ is extremal in $C_{[0, \infty)}(\mathcal{A})$. All that remains is to show extremality of functions (23) 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

The proof of Proposition 6.3 can be simplified considerably if we could start with the fact that

$$\Lambda_{[0, \infty)}^*(\mathcal{A}) = \{\lambda : \lambda_2 = -\lambda_1 = 1, \lambda_k = 0 \forall k \geq 3\} \cup \Lambda_{\mathbb{R}}^*(\mathcal{A}). \quad (24)$$

It is possible to prove the above identity without speaking of signomials, however the authors felt such an approach would obfuscate the primal decomposition-approach which is the focus of this section.

From a computational perspective it is easy to verify Equation 24. One may begin by defining $T_\beta = \{\phi_\lambda : \lambda \in \Lambda_X(\mathcal{A}, \beta)\} \cup \{(\mathbf{0}, 1)\}$, and then note that $G_X(\mathcal{A}) = \text{co} \left(\bigcup_{\beta \in \mathcal{A}} T_\beta \right)$. From here one may use Theorem 3.6 to find $\text{co} T_\beta = \{(\nu, t) \in N_\beta \times \mathbb{R} : \sigma_X(-\mathcal{A}\nu) \leq t\}$. In the case $X = [0, \infty)$, it is straightforward to compute the extreme rays of these $\text{co} T_\beta$ using a computational geometry toolbox (such as MPT3), and subsequently find the extreme rays of the polyhedral cone $G_X(\mathcal{A})$.

More generally, when X is a polyhedron, one may compute extreme rays of $G_X(\mathcal{A})$ by representing the polyhedral cones $\text{co} T_\beta$ as projections of higher-dimensional cones obtained by duality. For example, if $X = \{x : Ax + b \geq 0\}$, then

$$\text{co} T_\beta = \{(\nu, t) \in N_\beta \times \mathbb{R} : \exists \eta \geq 0 \text{ where } \mathcal{A}\nu = A^T \eta, b^T \eta \leq t\}.$$

To find the extreme rays for this set, one should apply Fourier-Motzkin elimination to project the lifted representation down to $\mathbb{R}^A \times \mathbb{R}$, and then proceed as usual.

⁵specifically when it holds with equality; also using $\exp(\alpha_i x) = (\exp(\alpha_{i+1}x))^{\lambda_{i+1}} (\exp(\alpha_{i-1}x))^{\lambda_{i-1}}$

7. DISCUSSION AND CONCLUSION

Over the course of this article we have introduced a new convex-geometric notion of a “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 (\mathcal{A}, X) -circuit allows an alternative construction of X -SAGE cones (Theorems 4.6 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 a 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. In closing, besides the natural questions of further understanding the X -SAGE cones, we offer a few more outreaching lines for follow-up work related to our results.

Of particular note is the task of formally situating our notion of (\mathcal{A}, X) -circuits in the context of matroid theory (in the case when X is a polyhedron). Here one can use an interpretation suggested by Theorem 3.7, that circuits $\lambda \in \Lambda_X(\mathcal{A}, \beta)$ correspond to facets of $-\mathcal{A}^T X + N_\beta^\circ$. There are also several ways to generalize our notion of an (\mathcal{A}, X) -circuit. For example, rather than requiring local strict-sublinearity of an augmented support function $\nu \mapsto \sigma_X(-\mathcal{A}\nu)$ confined to N_β , one could ask for local strict-sublinearity when restricted to some other closed and pointed cone $K \neq N_\beta$. Such an approach can be used to define “nonsimplicial” (\mathcal{A}, X) -circuits.

A broad area of follow-up work is in-depth analysis of multiplicatively-convex sets in Euclidean space (i.e. those $S \subset \mathbb{R}_+^n$ for which $\log(S) = \{t : \exp t \in S\}$ is a convex set). Some properties of this class of sets include closure under intersection, and closure under the induced-cone operation. Generally speaking, multiplicatively convex sets are most interesting when the logarithmic transform is invertible, i.e. when $S = \text{cl}(S \cap \mathbb{R}_{++}^n)$. Lemma 5.12 provided a result on separation theory for multiplicatively-convex sets with a certain structure; it would be valuable to see how the lemma’s construction generalizes to other multiplicatively-convex sets.

In terms of classical convex analysis and convex optimization, it is of interest to develop converse statements to Corollaries 4.7 and 4.8. For Corollary 4.7 one could try to show that for $X = \{x : \|x\|_2 \leq 1\}$, there exist choices of \mathcal{A} where $C_X(\mathcal{A})$ is not power-cone representable. For Corollary 4.8 one could take $X = \mathbb{R}^n$ and $\mathcal{A} = \{\beta, \delta_1, \dots, \delta_n\}$ where β is an irrational convex combination of the standard basis vectors $\{\delta_1, \dots, \delta_n\}$, and try to show that $C_X(\mathcal{A})$ does not admit a second-order-cone representation. Proving either of these converse statements can be seen as stepping stones towards conditions for which $C_X(\mathcal{A})$ is not semidefinite representable.

REFERENCES

- [1] B. Reznick. Forms derived from the arithmetic-geometric inequality. *Math. Annalen*, 283(3):431–464, 1989.
- [2] 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, May 2012.
- [3] S. Ilman and T. de Wolff. Amoebas, nonnegative polynomials and sums of squares supported on circuits. *Res. Math. Sci.*, 3(paper no. 9), 2016.
- [4] V. Chandrasekaran and P. Shah. Relative entropy relaxations for signomial optimization. *SIAM J. Optim.*, 26(2):1147–1173, 2016.

- [5] R. Murray, V. Chandrasekaran, and A. Wierman. Newton polytopes and relative entropy optimization. Preprint, arXiv:1810.01614, 2018.
- [6] J. Wang. Nonnegative polynomials and circuit polynomials. Preprint, arXiv:1804.09455, 2018.
- [7] L. Katthän, H. Naumann, and T. Theobald. A unified framework of SAGE and SONC polynomials and its duality theory, 2019. Preprint, arXiv:1903.08966.
- [8] J. Forsgård and T. de Wolff. The algebraic boundary of the SONC cone. Preprint, arXiv:1905.04776, 2019.
- [9] 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.
- [10] M. Dressler, S. Ilman, and T. De Wolff. An approach to constrained polynomial optimization via nonnegative circuit polynomials and geometric programming. *J. Symb. Comp.*, 91:149–172, 2019.
- [11] M. Dressler, A. Kurpisz, and T. de Wolff. Optimization over the Boolean hypercube via sums of nonnegative circuit polynomials. Preprint, arXiv:1802.10004, 2018.
- [12] 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.
- [13] R. Murray, V. Chandrasekaran, and A. Wierman. Signomial and polynomial optimization via relative entropy and partial dualization. Preprint, arXiv:1907.00814, 2019.
- [14] A. H. Wang, P. Jaini, Y. Yu, and P. Poupert. Complete hierarchy of relaxation for constrained signomial positivity. Preprint, arXiv:2003.03731, 2020.
- [15] A. Agrawal, S. Diamond, and S. Boyd. Disciplined geometric programming. *Optim. Letters*, 13(5):961–976, March 2019.
- [16] M. Özdemir, Ç. Yildiz, and M. Gürbüz. A note on geometrically convex functions. *J. Inequalities and Applications*, 2014(1):180, 2014.
- [17] W. Jarczyk and J. Matkowski. On Mulholland’s inequality. *Proc. Amer. Math. Soc.*, 130(11):3243–3247, 2002.
- [18] C. P. Niculescu. Convexity according to the geometric mean. *Math. Inequalities & Applications*, 3:155–167, 2000.
- [19] M. Dressler, J. Heuer, H. Naumann, and T. de Wolff. Global optimization via the dual SONC cone and linear programming. Preprint, arXiv:2002.09368, 2020.
- [20] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1997.
- [21] G. M. Ziegler. *Lectures on Polytopes*. Springer, New York, 1995.
- [22] S. Lu and S. M. Robinson. Normal fans of polyhedral convex sets. *Set-Valued Analysis*, 16(2-3):281–305, 2008.
- [23] D. Papp. Duality of sum of nonnegative circuit polynomials and optimal SONC bounds. Preprint, arXiv:1912.04718, 2019.
- [24] 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.
- [25] J. Wang and V. Magron. A second order cone characterization for sums of nonnegative circuits. Preprint, arXiv:1906.06179, 2019.
- [26] H. Naumann and T. Theobald. The \mathcal{S} -cone and a primal-dual view on second-order representability. Preprint, arXiv:2003.09495, 2020.
- [27] A. Ben-Tal and A. Nemirovski. *Lectures on Modern Convex Optimization*. SIAM, Philadelphia, 2001.

8. APPENDIX

We provide two auxiliary propositions regarding convex analysis, which are used in the proofs of Theorem 4.6 and Lemma 5.10.

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 $(c_1, \dots, c_m) \geq 0$, we have*

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

– 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 $D(\nu, c_{\setminus 0}) - 1^T \nu \leq c_0$ 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 $S \subset \mathbb{R}^m \setminus \{0\}$ 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 $t \in (\text{cl } T) \setminus \{0\}$. 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 [20, 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

RILEY MURRAY: CALIFORNIA INSTITUTE OF TECHNOLOGY (CALTECH) - DEPARTMENT OF COMPUTING & MATHEMATICAL SCIENCES, 1200 EAST CALIFORNIA BLVD, PASADENA, CA 91125, USA

HELEN NAUMANN, THORSTEN THEOBALD: GOETHE-UNIVERSITÄT, FB 12 – INSTITUT FÜR MATHEMATIK, POSTFACH 11 19 32, D-60054 FRANKFURT AM MAIN, GERMANY