



HAL
open science

Exploiting Sparsity for Semi-Algebraic Set Volume Computation

Matteo Tacchi, Tillmann Weisser, Jean-Bernard Lasserre, Didier Henrion

► **To cite this version:**

Matteo Tacchi, Tillmann Weisser, Jean-Bernard Lasserre, Didier Henrion. Exploiting Sparsity for Semi-Algebraic Set Volume Computation. *Foundations of Computational Mathematics*, 2021, 22, pp.161–209. 10.1007/s10208-021-09508-w . hal-02010175v2

HAL Id: hal-02010175

<https://hal.science/hal-02010175v2>

Submitted on 13 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Exploiting Sparsity for Semi-Algebraic Set Volume Computation*

M. TACCHI^{1,2} T. WEISSER³ J. B. LASSERRE^{1,4} D. HENRION^{1,5}

July 13, 2020

Abstract

We provide a systematic deterministic numerical scheme to approximate the volume (i.e. the Lebesgue measure) of a basic semi-algebraic set whose description follows a correlative sparsity pattern. As in previous works (without sparsity), the underlying strategy is to consider an infinite-dimensional linear program on measures whose optimal value is the volume of the set. This is a particular instance of a generalized moment problem which in turn can be approximated as closely as desired by solving a hierarchy of semidefinite relaxations of increasing size. The novelty with respect to previous work is that by exploiting the sparsity pattern we can provide a sparse formulation for which the associated semidefinite relaxations are of much smaller size. In addition, we can decompose the sparse relaxations into completely decoupled subproblems of smaller size, and in some cases computations can be done in parallel. To the best of our knowledge, it is the first contribution that exploits correlative sparsity for volume computation of semi-algebraic sets which are possibly high-dimensional and/or non-convex and/or non-connected.

1 Introduction

This paper is in the line of research concerned with computing approximations of the volume (i.e. the Lebesgue measure) of a given compact basic semi-algebraic set \mathbf{K} of \mathbb{R}^n neither necessarily convex nor connected. Computing or even approximating the volume of a convex body is hard theoretically and in practice as well. Even if \mathbf{K} is a convex polytope, exact computation of its volume or integration over \mathbf{K} is a computational challenge. Computational complexity of these problems is discussed in, e.g. [1, 2, 3]. In particular, any deterministic algorithm with polynomial-time complexity that would compute an upper bound and a lower bound on the volume cannot yield an estimate on the bound ratio better than polynomial in the dimension n . For more detail, the interested reader is referred to the discussion in [4] and to [5] for a comparison.

If one accepts randomized algorithms that fail with small probability, then the situation is more favorable. Indeed, the probabilistic approximation algorithm of [6] computes the volume to fixed arbitrary relative precision $\epsilon > 0$ in time polynomial in $1/\epsilon$. The algorithm uses approximation schemes based on rapidly mixing Markov chains and isoperimetric inequalities, see also hit-and-run algorithms described in e.g. [7, 8, 9]. So far, it seems that the recent work [10] has provided the best algorithm of this type.

¹LAAS-CNRS, Université de Toulouse, France.

²Réseau de Transport d'Electricité (RTE), France.

³Theoretical Division (T-5/CNLS), Los Alamos National Laboratory, Los Alamos, USA.

⁴Institut de Mathématiques de Toulouse, Université de Toulouse, France.

⁵Faculty of Electrical Engineering, Czech Technical University in Prague, Czechia.

*This work was partly funded by the RTE company and by the ERC Advanced Grant Taming.

The moment approach for volume computation

In [4] a general deterministic methodology was proposed for approximating the volume of a compact basic semi-algebraic set \mathbf{K} , not necessarily convex or connected. It was another illustration of the versatility of the so-called *moment-SOS (sums of squares) hierarchy* developed in [11] for solving the Generalized Moment Problem (GMP) with polynomial data.

Briefly, the underlying idea is to view the volume as the optimal value of a GMP, i.e., an infinite-dimensional Linear Program (LP) on an appropriate space of finite Borel measures. Then one may approximate the value from above by solving a hierarchy of *semidefinite programming (SDP) relaxations* with associated sequence of optimal values indexed by an integer d . Monotone convergence of the bounds is guaranteed when d increases. Extensions to more general measures and possibly non-compact sets were then proposed in [12]. The order d in the hierarchy encodes the amount of information that is used, namely the number of moments of the Lebesgue measure up to degree d . It is a crucial factor for the size of the associated SDP problem. More precisely, the size grows in d^n as d increases, which so far limits its application to sets of small dimension n , typically up to 4 or 5.

In view of the growth of the size of the SDP problem with increasing order it is desirable to converge quickly towards the optimal value of the LP. However, this convergence is expected to be slow in general. One reason becomes clear when looking at the dual LP which attempts to compute a polynomial approximation (from above) of the (discontinuous) indicator function of the set \mathbf{K} . Hence one is then faced with a typical Gibbs effect¹, well-known in the theory of approximation. To overcome this drawback the authors in [4] have proposed to use an alternative criterion for the LP, which results in a significantly faster convergence. However in doing so, the monotonicity of the convergence (a highly desirable feature to obtain a sequence of upper bounds) is lost. In the related work [12] the author has proposed an alternative strategy which consists of strengthening the relaxations by incorporating additional linear moment constraints known to be satisfied at the optimal solution of the LP. These constraints come from a specific application of Stokes' theorem. Remarkably, adding these *Stokes constraints* results in a significantly faster convergence while keeping monotonicity.

Motivation

The measure approach for the volume computation problem is intimately linked to the use of occupation measures, in dynamical systems theory, for computing the region of attraction (ROA) of a given target set. Indeed, in [13], the problem of estimating the ROA is formulated as a GMP very similar to the volume computation problem. The idea is to maximize the volume of a set of initial conditions that yield trajectories ending in the target set after a given time.

This problem of estimating the ROA of a target set is crucial in power systems safety assessment, since the power grids must have good stability properties. Perturbations (such as short-circuits or unscheduled commutations) should be tackled before they get the system outside the ROA of the nominal operating domain. Currently the stability of power grids is estimated through numerous trajectory simulations that prove very expensive. In the wake of the energy transition, it is necessary to find new tools for estimating the stability of power systems. The conservative, geometric characterization of the region of attraction as formulated in [13] is a very promising approach for this domain of application, see [14].

As in volume computation, the main limitation of this method is that only problems of modest dimension can be handled by current solvers. Exploiting sparsity seems to be the best approach to allow scalability both in volume computation and ROA estimation. Since volume estimation is a simpler instance of the GMP than ROA estimation, we decided to address first the former problem.

¹The Gibbs effect appears at a jump discontinuity when one approximates a piecewise continuously differentiable function with a continuous function, e.g. by its Fourier series.

In addition, volume computation with respect to a measure satisfying some conditions (e.g. compactly supported or Gaussian measure) also has applications in the fields of geometry and probability computation, which is the reason why many algorithms were already proposed for volume computation of convex polytopes and convex bodies in general.

Contribution

We design deterministic methods that provide approximations with strong asymptotic guarantees of convergence to the volume of \mathbf{K} . The methodology that we propose is similar in spirit to the one initially developed in [4] as described above and its extension to non-compact sets and Gaussian measures of [12]. However it is not a straightforward or direct extension of [4] or [12], and it has the following important distinguishing features:

(i) It can handle sets $\mathbf{K} \subset \mathbb{R}^n$ of potentially large dimension n provided that some sparsity pattern (namely: correlative sparsity, see section 2.3 as well as [15, 16] for details) is present in the description of \mathbf{K} . This is in sharp contrast with [4].

(ii) The computation of upper and lower bounds can be decomposed into smaller independent problems of the same type, and depending on the sparsity pattern, some of the computations can even be done in parallel. This fact alone is remarkable and unexpected.

To the best of our knowledge, this is the first deterministic method for volume computation that takes benefit from a correlative sparsity pattern in the description of \mathbf{K} in the two directions of (a) decomposition into problems of smaller size and (b) parallel computation. Of course this sharp improvement is performed at some price: our framework only works on semi-algebraic sets that present the appropriate correlative sparsity pattern (see Assumption 16 as well as section 6.5 and appendix A for detailed discussion on its applicability).

The key idea is to provide a new and very specific *sparse formulation* of the original problem in which one defines a set of marginal measures whose (small dimensional) support is in accordance with the correlative sparsity pattern present in the description of the set \mathbf{K} . However, those marginal measures are not similar to the ones used in the sparse formulation [17] of polynomial optimization problems over the same set \mathbf{K} . Indeed they are not expected to satisfy the consistency condition of [17]².

Finally, in principle, our floating point volume computation in large dimension n is faced with a crucial numerical issue. Indeed as in Monte-Carlo methods, up to rescaling, one has to include the set \mathbf{K} into a box \mathbf{B} of unit volume. Therefore the volume of \mathbf{K} is of the order ϵ^n for some $\epsilon \in (0, 1)$ and therefore far beyond machine precision as soon as n is large. To handle this critical issue we develop a *sparsity-adapted rescaling* which allows us to compute very small volumes in potentially very high dimension with good precision.

A motivating example

Consider the following set

$$\mathbf{K} := \{\mathbf{x} \in [0, 1]^{100} : x_i x_{i+1} \leq 1/2, i = 1, \dots, 99\}.$$

This is a *high-dimensional non-convex sparse semi-algebraic set*. The precise definition of a sparse semi-algebraic set will be given later on, but so far notice that in the description of \mathbf{K} each constraint involves only 2 variables out of 100. The volume of \mathbf{K} is hard to compute, but thanks to the structured description of the set we are able to prove numerically that its volume is smaller than $2 \cdot 10^{-5}$ in less than 2 minutes on a standard computer.

For this we have implemented a specific version of the moment-SOS hierarchy of SDP relaxations to solve the GMP, in which we exploit the correlative sparsity pattern of the set \mathbf{K} . The

²If two measures share some variables then the consistency condition requires that their respective marginals coincide.

basic idea is to replace the original GMP that involves an unknown measure on \mathbb{R}^{100} (whose SDP relaxations are hence untractable) with a GMP involving 99 measures on \mathbb{R}^2 (hence tractable). In addition, this new GMP can be solved either in one shot (with the 99 unknown measures) or by solving sequentially 99 GMPs involving (i) one measure on \mathbb{R}^2 and (ii) some data obtained from the GMP solved at previous step. Our approach can be sketched as follows.

First, we rescale the problem so that the set \mathbf{K} is included in the unit box $\mathbf{B} := [0, 1]^n$ on which the moments of the Lebesgue measure are easily computed.

Next, we describe the volume problem on \mathbf{K} as a chain of volume subproblems on the subspaces $\text{Im}(\pi_i)$ where $\pi_i(x_1, \dots, x_{100}) = (x_i, x_{i+1})$, with a link between the i -th and $(i+1)$ -th sub-problems.

Finally, in this example, as $n = 100$ and $\mathbf{K} \subset \mathbf{B}$, the volume of \mathbf{K} is very small and far below standard floating point machine precision. To handle this numerical issue, we have implemented a sparsity-adapted strategy which consists of rescaling each subproblem defined on the projections of \mathbf{K} to obtain intermediate values all with the same order of magnitude. Once all computations (involving quantities of the same order of magnitude) have been performed, the correct value of the volume is obtained by a reverse scaling.

The sparse formulation stems from considering some measure marginals appropriately defined according to the correlative sparsity pattern present in the description of \mathbf{K} . It leads to a variety of algorithms to compute the volume of sparse semi-algebraic sets.

The outline of the paper is as follows. In Section 2 we describe briefly the moment-SOS hierarchy for semi-algebraic set volume approximation, as well as our notion of sparse semi-algebraic set. In Section 3, we introduce a first constructive theorem that allows to efficiently compute the volume of specific sparse sets with the hierarchy. Section 4 is dedicated to a method for accelerating the convergence of the sparse hierarchy, and in section 5 we discuss some numerical examples. Eventually, Section 6 presents the general theorem for computing the volume of a more general variety of sparse sets using parallel computations, accompanied with some other examples.

2 Preliminaries

2.1 Notations and definitions

Given a closed set $\mathbf{K} \subset \mathbb{R}^n$, we denote by $\mathcal{C}(\mathbf{K})$ the space of continuous functions on \mathbf{K} , $\mathcal{M}(\mathbf{K})$ the space of finite signed Borel measures on \mathbf{K} , with $\mathcal{C}_+(\mathbf{K})$ and $\mathcal{M}_+(\mathbf{K})$ their respective cones of positive elements.

The Lebesgue measure on \mathbb{R}^n is $\lambda^n(d\mathbf{x}) := dx_1 \otimes dx_2 \otimes \dots \otimes dx_n = d\mathbf{x}$, and its restriction to a set $\mathbf{K} \subset \mathbb{R}^n$ is $\lambda_{\mathbf{K}} := \mathbb{1}_{\mathbf{K}} \lambda^n$, where $\mathbb{1}_{\mathbf{K}}$ denotes the indicator function equal to 1 in \mathbf{K} and zero outside. In this paper, we focus on computing the volume or Lebesgue measure of \mathbf{K} , that we denote by $\text{vol } \mathbf{K}$ or $\lambda^n(\mathbf{K})$ or $\lambda_{\mathbf{K}}(\mathbb{R}^n)$.

Given a Euclidean space \mathbf{X} and a subspace $\mathbf{X}_i \subset \mathbf{X}$, the orthogonal projection map from \mathbf{X} to \mathbf{X}_i is denoted by $\pi_{\mathbf{X}_i}$. Let \mathbf{e}_j denote the j -th vector of the canonical basis of \mathbb{R}^n such that if $\mathbf{x} = (x_1, \dots, x_n)$ then $x_j = \mathbf{x} \cdot \mathbf{e}_j$, where the dot denotes the scalar product between vectors. The m -dimensional subspace spanned by vectors $\mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_m}$ is denoted $\langle x_{i_1}, \dots, x_{i_m} \rangle$ or $\langle x_{i_j} \rangle_{1 \leq j \leq m}$. Given a measure $\mu \in \mathcal{M}(\mathbf{X})$, its marginal with respect to \mathbf{X}_i is denoted by $\mu^{\mathbf{X}_i} \in \mathcal{M}(\mathbf{X}_i)$. It is equal to the image or push-forward measure of μ through the map $\pi_{\mathbf{X}_i}$.

Let $\mathbb{R}[\mathbf{x}]$ be the ring of polynomials in the variables $\mathbf{x} = (x_1, \dots, x_n)$ and let $\mathbb{R}[\mathbf{x}]_d$ be the vector space of polynomials of degree at most d , whose dimension is $s(d) := \binom{n+d}{n}$. For every $d \in \mathbb{N}$, let $\mathbb{N}_d^n := \{\alpha \in \mathbb{N}^n : |\alpha| (= \sum_i \alpha_i) \leq d\}$, and let $\mathbf{v}_d(\mathbf{x}) = (\mathbf{x}^\alpha)_{\alpha \in \mathbb{N}_d^n}$ be the vector of monomials of the canonical basis $(\mathbf{x}^\alpha)_{\alpha \in \mathbb{N}_d^n}$ of $\mathbb{R}[\mathbf{x}]_d$. A polynomial $p \in \mathbb{R}[\mathbf{x}]_d$ is written as

$$\mathbf{x} \mapsto p(\mathbf{x}) = \sum_{\alpha \in \mathbb{N}^n} p_\alpha \mathbf{x}^\alpha = \mathbf{p} \cdot \mathbf{v}_d(\mathbf{x})$$

for some vector of coefficients $\mathbf{p} = (p_\alpha)_{\alpha \in \mathbb{N}_d^n} \in \mathbb{R}^{s(d)}$.

We say that $\mu \in \mathcal{M}(\mathbb{R}^n)$ is a representing measure of the sequence $\mathbf{m} = (m_\alpha)_{\alpha \in \mathbb{N}^n} \subset \mathbb{R}$ whenever

$$m_\alpha = \int \mathbf{x}^\alpha d\mu(\mathbf{x})$$

for all $\alpha \in \mathbb{N}^n$. Given a sequence $\mathbf{m} = (m_\alpha)_{\alpha \in \mathbb{N}^n}$, let $L_{\mathbf{m}} : \mathbb{R}[\mathbf{x}] \rightarrow \mathbb{R}$ be the linear functional

$$f \left(= \sum_{\alpha} f_{\alpha} \mathbf{x}^{\alpha} \right) \mapsto L_{\mathbf{m}}(f) := \sum_{\alpha} f_{\alpha} m_{\alpha}.$$

Given a sequence $\mathbf{m} = (m_\alpha)_{\alpha \in \mathbb{N}^n}$, and a polynomial $g := \sum_{\gamma} g_{\gamma} \mathbf{x}^{\gamma} \in \mathbb{R}[\mathbf{x}]$, the localizing moment matrix of order d associated with \mathbf{m} and g is the real symmetric matrix $\mathbf{M}_d(g, \mathbf{m})$ of size $s(d)$ with rows and columns indexed in \mathbb{N}_d^n and with entries

$$\begin{aligned} \mathbf{M}_d(g, \mathbf{m})(\alpha, \beta) &:= L_{\mathbf{m}}(g(\mathbf{x}) \mathbf{x}^{\alpha+\beta}) \\ &= \sum_{\gamma} g_{\gamma} m_{\alpha+\beta+\gamma}, \quad \alpha, \beta \in \mathbb{N}_d^n. \end{aligned}$$

When $g \equiv 1$, the localizing moment matrix $\mathbf{M}_d(\mathbf{m})$ is called simply the moment matrix.

2.2 The Moment-SOS hierarchy for volume computation

Let $\mathbf{B} := [0, 1]^n \subset \mathbf{X} := \mathbb{R}^n$ be the n -dimensional unit box, and let $\mathbf{K} \subset \mathbf{B}$ be a closed basic semialgebraic set defined by

$$\mathbf{K} := \{\mathbf{x} \in \mathbf{X} : g_i(\mathbf{x}) \geq 0, i = 1, \dots, m\} = \{\mathbf{x} \in \mathbf{X} : \mathbf{g}(\mathbf{x}) \geq 0\}$$

where $\mathbf{g} = (g_i)_{i=1, \dots, m} \in \mathbb{R}[\mathbf{x}]^m$ and the rightmost vector inequality is meant entrywise. As in [4] consider the infinite-dimensional linear program (LP) on measures

$$\begin{aligned} \max_{\mu, \hat{\mu} \in \mathcal{M}_+(\mathbf{B})} & \int d\mu \\ \text{s.t.} & \mu + \hat{\mu} = \lambda_{\mathbf{B}} \\ & \text{spt } \mu \subset \mathbf{K} \\ & \text{spt } \hat{\mu} \subset \mathbf{B}. \end{aligned} \tag{1}$$

Its value is equal to $\text{vol } \mathbf{K}$ and the measures $\mu^* = \lambda_{\mathbf{K}}$, $\hat{\mu}^* := \lambda_{\mathbf{B} \setminus \mathbf{K}}$ are the unique optimal solutions of (1). The dual of (1) is the infinite-dimensional LP on continuous functions

$$\begin{aligned} \inf_{v \in \mathcal{C}_+(\mathbf{B})} & \int v d\lambda_{\mathbf{B}} \\ \text{s.t.} & v \geq \mathbb{1}_{\mathbf{K}}. \end{aligned} \tag{2}$$

It turns out that there is no duality gap between (1) and (2), i.e., they both have the same optimal value. Notice that a minimizing sequence of (2) approximates the indicator function $\mathbb{1}_{\mathbf{K}}$ from above by polynomials of increasing degrees.

The LP (1) is a particular and simple instance of the Generalized Moment Problem (GMP). As described in [4] one may approximate its optimal value as closely as desired by using the following key result in [11]. Given an infinite dimensional LP on measures, one can construct a hierarchy of finite dimensional semidefinite programs³ (SDP) whose associated sequence of optimal values converges monotonically to the optimal value of the original infinite dimensional LP. The basic idea is to represent a measure with the sequence \mathbf{m} of its moments, and to

³A semidefinite program is a convex conic optimization problem that can be solved numerically efficiently, e.g. by using interior point methods.

formulate finite dimensional SDPs on truncations of the sequence \mathbf{m} . When this strategy is applied to LP (1), the step d of the moment-SOS hierarchy consists of solving the SDP relaxation

$$\begin{aligned} \mathbf{P}_d : \quad & \max_{\mathbf{m}, \widehat{\mathbf{m}} \in \mathbb{R}^{s(d)}} m_0 \\ & \text{s.t.} \quad m_\alpha + \widehat{m}_\alpha = \int_{\mathbf{B}} \mathbf{x}^\alpha d\mathbf{x}, \quad \alpha \in \mathbb{N}_{2d}^n \\ & \quad \mathbf{M}_d(\mathbf{m}) \succeq 0, \mathbf{M}_d(\widehat{\mathbf{m}}) \succeq 0 \\ & \quad \mathbf{M}_{d-d_i}(g_i \mathbf{m}) \succeq 0, \quad i = 1, \dots, m \end{aligned} \quad (3)$$

where $\mathbf{m} = (m_\alpha)_{\alpha \in \mathbb{N}_{2d}^n}$, $\widehat{\mathbf{m}} = (\widehat{m}_\alpha)_{\alpha \in \mathbb{N}_{2d}^n}$, and $d_i = \lceil (\deg g_i)/2 \rceil$, $i = 1, \dots, m$.

The sequence of SDP problems $(\mathbf{P}_d)_{d \in \mathbb{N}}$ indexed by the relaxation order d is a hierarchy in the sense that its sequence of values converges monotonically from above to $\text{vol } \mathbf{K}$ when d increases. Each SDP relaxation \mathbf{P}_d has a dual formulated in terms of sums of squares (SOS) of polynomials, which leads to a dual SOS hierarchy, whence the name *moment-SOS hierarchy*. The basic moment-SOS hierarchy can be modeled using the GloptiPoly Matlab toolbox [18] and solved using any SDP solver, e.g. SeDuMi or Mosek. For more details on the moment-SOS hierarchy and some of its applications, the interested reader is referred to [11].

2.3 The correlative sparsity pattern and its graph representation

This work heavily relies on a specific notion of sparsity defined as follows.

Definition 1. A scalar polynomial p is said to be *sparse* when its vector of coefficients \mathbf{p} is sparse. In other words, p is a linear combination of a small number of monomials.

Definition 2. A family of polynomial vectors $(\mathbf{g}_1, \dots, \mathbf{g}_m)$ is said to be *correlatively sparse* whenever its correlative sparsity pattern matrix $\mathbf{R} := (R_{ij})_{1 \leq i, j \leq m}$, defined by

$$R_{ij} := \delta_{ij} + \sum_{k=1}^m \left\| \frac{\partial}{\partial x_i} \mathbf{g}_k \right\| \left\| \frac{\partial}{\partial x_j} \mathbf{g}_k \right\|$$

(where $\delta_{ij} = 1$ if $i = j$ and 0 otherwise, and $\|\cdot\|$ is any norm on polynomial vectors), is sparse. In other words, for many pairs of indices $i \neq j$, the variables x_i and x_j do not appear simultaneously in any element of $\{\mathbf{g}_1, \dots, \mathbf{g}_m\}$.

Definition 3. The *correlation graph* $G = (V, E)$ of $(\mathbf{g}_1, \dots, \mathbf{g}_m)$ is defined by vertices $V = \{1, \dots, m\}$ and edges $E = \{(i, j) \in \{1, \dots, m\}^2 : i \neq j \ \& \ R_{ij} \neq 0\}$. The *correlative sparsity CS* of $(\mathbf{g}_1, \dots, \mathbf{g}_m)$ is the treewidth of its correlation graph.

This paper proposes a method to reduce the size of the volume computation SDP to $\binom{CS+d+1}{d}$ instead of $\binom{n+d}{d}$, under appropriate assumptions.

Definition 4. The *support* of \mathbf{g}_i is the set $S(\mathbf{g}_i) := \{j \in \{1, \dots, n\} : \frac{\partial}{\partial x_j} \mathbf{g}_i \neq 0\}$. The *support subspace* of \mathbf{g}_i is the set $\mathbf{X}_i := \langle x_j \rangle_{j \in S(\mathbf{g}_i)}$, whose dimension is smaller than n . Since by definition $\mathbf{g}_i = \mathbf{g}_i \circ \pi_{\mathbf{X}_i}$, we use both notations with the same meaning. Then $\mathbf{X} := \sum_{i=1}^m \mathbf{X}_i$ is called the *coordinate subspace decomposition* associated to $(\mathbf{g}_1, \dots, \mathbf{g}_m)$.

Without loss of generality, we can suppose that $\mathbf{X} = \mathbb{R}^n$ (otherwise, there would be variables that appear in none of the \mathbf{g}_i s).

Definition 5. A *sparse basic semi-algebraic set* has a description

$$\mathbf{K} := \{\mathbf{x} \in \mathbf{X} : \mathbf{g}_i(\pi_{\mathbf{X}_i}(\mathbf{x})) \geq 0, \ i = 1, \dots, m\}$$

where $(\mathbf{g}_i)_{1 \leq i \leq m}$ is a correlatively sparse family of polynomial vectors (inequalities are meant entrywise) and $\mathbf{X} = \sum_{i=1}^m \mathbf{X}_i$ is the coordinate subspace decomposition associated to $(\mathbf{g}_i)_{1 \leq i \leq m}$ (and, by extension, to \mathbf{K}). A *sparse semi-algebraic set* is a finite union of sparse basic semi-algebraic sets that share the same coordinate subspace decomposition.

A simple example of a sparse basic semi-algebraic set is

$$\mathbf{K} := \{\mathbf{x} = (x_1, x_2, x_3, x_4) \in \mathbb{R}^4 : \mathbf{g}_1(x_1, x_2) \geq 0, \mathbf{g}_2(x_2, x_3) \geq 0, \mathbf{g}_3(x_3, x_4) \geq 0\} \quad (4)$$

for $\mathbf{X} = \mathbb{R}^4$, $\mathbf{X}_1 = \langle x_1, x_2 \rangle$, $\mathbf{X}_2 = \langle x_2, x_3 \rangle$, $\mathbf{X}_3 = \langle x_3, x_4 \rangle$ and the projection maps are $\pi_{\mathbf{X}_1}(\mathbf{x}) = (x_1, x_2)$, $\pi_{\mathbf{X}_2}(\mathbf{x}) = (x_2, x_3)$, $\pi_{\mathbf{X}_3}(\mathbf{x}) = (x_3, x_4)$.

Our methodology is based on the classical theory of clique trees. Up to a chordal extension (which is equivalent to slightly weakening the correlative sparsity pattern), we suppose that the correlation graph is *chordal* (i.e. every cycle of length greater than 3 has a chord, i.e. an edge linking two nonconsecutive vertices). Then we construct the maximal *cliques* of the graph (a clique C is a subset of vertices such that every vertex of C is connected to all the other vertices of C , a clique is maximal when its cardinal is maximal). Most of the time, up to concatenation of some of the \mathbf{g}_i the maximal cliques of a chordal correlation graph are exactly the supports of the \mathbf{g}_i : $C_i = S(\mathbf{g}_i)$, so in the following we will consider only such case⁴. Figure 1 illustrates this construction for the sparse set (4), the vertices are denoted by x_i and our maximal cliques are denoted by C_j .




Figure 1: Graph associated to the sparse set (4).

Then, we construct a *clique tree* which is instrumental to the computer implementation. It is proved in [19] that if the graph is chordal, then its maximal cliques can be organized within a tree satisfying the clique intersection property: for two maximal cliques C and C' the intersection $C \cap C'$ is contained in every maximal clique on the path from C to C' . Figure 2 represents the clique tree associated to the sparse set (4).




Figure 2: Linear clique tree associated to the sparse set (4).

For a slightly more complicated illustration, consider the sparse set

$$\mathbf{K} := \{\mathbf{x} \in \mathbb{R}^6 : \mathbf{g}_1(x_1, x_2) \geq 0, \mathbf{g}_2(x_2, x_3, x_4) \geq 0, \mathbf{g}_3(x_3, x_5) \geq 0, \mathbf{g}_4(x_4, x_6) \geq 0\} \quad (5)$$

whose correlation graph is represented on Figure 3 and whose clique tree is represented on Figure 4. The clique tree of Figure 2 is called *linear* because all maximal cliques form a single chain (i.e. they are in a sequence) with no branching. In contrast, the clique tree of Figure 4 is called *branched*.

Our method consists of conveniently rooting the clique tree and splitting the volume computation problem into lower-dimensional subproblems that are in correspondence with the maximal cliques of the graph. The subproblem associated with a maximal clique C takes as only input the solutions of the subproblems associated with the children of C in the clique tree. This way, one can compute in parallel the solutions of all the subproblems of a given generation, and then use their results to solve the subproblems of the parent generation. This is the meaning of the arrows in Figures 2 and 4. The volume of \mathbf{K} is the optimal value of the (last) sub-problem associated with the root C_1 of the tree.

⁴The only exception would be cliques forming a triangle and is tackled in detail in section 6.5.




Figure 3: Graph associated to the sparse set (5).




Figure 4: Branched clique tree associated to the sparse set (5).

3 Linear sparse volume computation

In this section we describe the method in the prototype case of linear clique trees. The more general case of branched clique trees is treated later on.

3.1 An illustrative example: the bicylinder

Before describing the methodology in the general case, we briefly explain the general underlying idea on a simple illustrative example. Consider the sparse semi-algebraic set

$$\mathbf{K} := \{\mathbf{x} \in \mathbb{R}^3 : \mathbf{g}_1(x_1, x_2) := 1 - x_1^2 - x_2^2 \geq 0, \mathbf{g}_2(x_2, x_3) := 1 - x_2^2 - x_3^2 \geq 0\} \quad (6)$$

modelling the intersection of two truncated cylinders $\mathbf{K}_1 := \{\mathbf{x} \in \mathbb{R}^3 : x_1^2 + x_2^2 \leq 1\}$ and $\mathbf{K}_2 := \{\mathbf{x} \in \mathbb{R}^3 : x_2^2 + x_3^2 \leq 1\}$, see Figure 3.1. The subspaces are $\mathbf{X}_1 = \langle x_1, x_2 \rangle$ and $\mathbf{X}_2 = \langle x_2, x_3 \rangle$ and the projection maps are $\pi_{\mathbf{X}_1}(\mathbf{x}) = (x_1, x_2)$ and $\pi_{\mathbf{X}_2}(\mathbf{x}) = (x_2, x_3)$. Let $\mathbf{U}_i := \pi_{\mathbf{X}_i}(\mathbf{K}_i)$ for $i = 1, 2$.

Following [4], computing $\text{vol} \mathbf{K}$ is equivalent to solving the infinite-dimensional LP (1). Next observe that in the description (6) of \mathbf{K} there is no direct interaction between variables x_1 and x_3 , but this is neither exploited in the LP formulation (1) nor in the SDP relaxations (3) to solve




Figure 5: A representation of the bicylinder produced by the AutoDesk Fusion 360 software.

(1). To exploit this correlative sparsity pattern we propose the following alternative formulation

$$\begin{aligned}
\text{vol } \mathbf{K} &= \max_{\substack{\mu_i \in \mathcal{M}_+(\mathbf{X}_i) \\ i=1,2}} \int_{\mathbb{R}^2} d\mu_1 & (7) \\
\text{s.t. } & \mu_2 \leq \lambda \otimes \lambda \\
& \mu_1 \leq \lambda \otimes \mu_2^{\langle x_2 \rangle} \\
& \text{spt } \mu_1 \subset \mathbf{U}_1, \text{ spt } \mu_2 \subset \mathbf{U}_2
\end{aligned}$$

where $\mu_2^{\langle x_2 \rangle}$ denotes the marginal of μ_2 in the variable x_2 .

In the sparse case, the basic idea behind our reformulation of the volume problem is as follows. We are interested in $\text{vol } \mathbf{K}$. However, as the marginal of a measure has the same mass as the measure itself, instead of looking for the full measure μ in problem (1), we only look for its marginal on \mathbf{X}_1 .

This marginal $\mu^{\mathbf{X}_1}$ is modeled by μ_1 in (7). In order to compute it, we need some additional information on μ captured by the measure μ_2 in (7). The unique optimal solution μ of (1) is

$$\begin{aligned}
d\mu(\mathbf{x}) &= d\lambda_{\mathbf{K}}(\mathbf{x}) \\
&= \mathbb{1}_{\mathbf{U}_1}(x_1, x_2) \mathbb{1}_{\mathbf{U}_2}(x_2, x_3) d\mathbf{x}
\end{aligned}$$

and therefore its marginal $\mu_1 := \mu^{\mathbf{X}_1}$ on (x_1, x_2) is

$$\begin{aligned}
d\mu_1(x_1, x_2) &= \int_0^1 d\mu(x_1, x_2, x_3) \\
&= \mathbb{1}_{\mathbf{U}_1}(x_1, x_2) dx_1 \underbrace{\left(\int_0^1 \mathbb{1}_{\mathbf{U}_2}(x_2, x_3) dx_3 \right)}_{d\mu_2^{\langle x_2 \rangle}(x_2)} dx_2 & (8)
\end{aligned}$$

where

$$d\mu_2(x_2, x_3) = d\lambda_{\mathbf{U}_2}(x_2, x_3). \quad (9)$$

What is the gain in solving (7) when compared to solving (1) ? Observe that in (7) we have two unknown measures μ_1 and μ_2 on \mathbb{R}^2 , instead of a single measure μ on \mathbb{R}^3 in (1). In the resulting SDP relaxations associated with (7) this translates into SDP constraints of potentially much smaller size. For instance, and to fix ideas, for the same relaxation degree d :

- The SDP relaxation \mathbf{P}_d associated with (1) contains a moment matrix (associated with μ in (1)) of size $\binom{3+d}{d}$;
- The SDP relaxation \mathbf{P}_d associated with (7) contains two moment matrices, one associated with μ_1 of size $\binom{2+d}{d}$, and one associated with μ_2 of size $\binom{2+d}{d}$, where μ_1 and μ_2 are as in (7).

As the size of those matrices is the crucial parameter for all SDP solvers, one can immediately appreciate the computational gain that can be expected from the formulation (7) versus the formulation (1) when the dimension is high or the relaxation order increases. Next it is not difficult to extrapolate that the gain can be even more impressive in the case where the correlative sparsity pattern is of the form

$$\mathbf{K} = \{(\mathbf{x}_0, \dots, \mathbf{x}_m) \in \mathbf{X} : g_i(\mathbf{x}_{i-1}, \mathbf{x}_i) \geq 0, \quad i = 1, \dots, m\}, \quad (10)$$

with $\mathbf{x}_i \in \mathbb{R}^{n_i}$ and $n_i \ll n$ for $i = 0, \dots, m$. In fact, it is straightforward to define examples of sets \mathbf{K} of the form (10) where the first SDP relaxation associated with the original dense LP formulation (1) cannot be even implemented on state-of-the-art computers, whereas the SDP relaxations associated with a generalization of the sparse LP formulation (7) can be easily implemented, at least for reasonable values of d .

3.2 Linear computation theorem

Let

$$\mathbf{K}_i := \{\mathbf{x} \in \mathbf{X} : \mathbf{g}_i(\pi_{\mathbf{X}_i}(\mathbf{x})) \geq 0\}$$

with $\mathbf{g}_i \in \mathbb{R}[\mathbf{x}_i]^{p_i}$, so that our sparse semi-algebraic set can be written

$$\mathbf{K} = \bigcap_{i=1}^m \mathbf{K}_i.$$

Moreover, let

$$\mathbf{U}_i := \{\mathbf{x}_i \in \mathbf{X}_i : \mathbf{g}_i(\mathbf{x}_i) \geq 0\} = \pi_{\mathbf{X}_i}(\mathbf{K}_i)$$

and let

$$\mathbf{Y}_i := \mathbf{X}_i \cap \mathbf{X}_{i+1}^\perp = \langle x_j \rangle_{j \in C_i \cap C_{i+1}^c}$$

be a subspace of dimension $n_i := |C_i \cap C_{i+1}^c|$ for $i = 1, \dots, m-1$ with $\mathbf{Y}_m = \mathbf{X}_m$. The superscript \perp denotes the orthogonal complement.

Assumption 6. For all $i \in \{2, \dots, m\}$ it holds $\mathbf{X}_i \cap \sum_{j=1}^{i-1} \mathbf{X}_j \neq \{0\}$.

Assumption 7. For all $i \in \{2, \dots, m\}$ it holds $\mathbf{X}_i \cap \sum_{j=1}^{i-1} \mathbf{X}_j \subset \mathbf{X}_{i-1}$.

If Assumption 6 is violated then \mathbf{K} can be decomposed as a Cartesian product, and one should just apply our methodology to each one of its factors. Assumption 7 ensures that the associated clique tree is linear.

Theorem 8. If Assumptions 6 and 7 hold, then $\text{vol } \mathbf{K}$ is the value of the LP problem

$$\max_{\substack{\mu_i \in \mathcal{M}_+(\mathbf{X}_i) \\ i=1, \dots, m}} \int d\mu_1 \quad (11)$$

$$s.t. \quad \mu_i \leq \mu_{i+1}^{\mathbf{X}_i \cap \mathbf{X}_{i+1}} \otimes \lambda^{n_i} \quad i = 1, \dots, m-1 \quad (12)$$

$$\mu_m \leq \lambda^{n_m} \quad (13)$$

$$\text{spt } \mu_i \subset \mathbf{U}_i \quad i = 1, \dots, m. \quad (14)$$

Proof. Let us first prove that the value of the LP is larger than $\text{vol } \mathbf{K}$. For $i = 1, \dots, m$, let $\mathbf{Z}_i := \mathbf{X}_i^\perp \cap \sum_{j=i+1}^m \mathbf{X}_j$ so that $\sum_{j=i}^m \mathbf{X}_j = \mathbf{X}_i \oplus \mathbf{Z}_i$. For $\mathbf{x}_i \in \mathbf{X}_i$ define

$$d\mu_i(\mathbf{x}_i) := \mathbb{1}_{\mathbf{U}_i}(\mathbf{x}_i) \left(\int_{\mathbf{Z}_i} \prod_{j=i+1}^m \mathbb{1}_{\mathbf{U}_j} \circ \pi_{\mathbf{X}_j}(\mathbf{x}_i + \mathbf{z}_i) d\mathbf{z}_i \right) d\mathbf{x}_i. \quad (15)$$

By construction $\mu_i \in \mathcal{M}_+(\mathbf{X}_i)$ and constraints (14) are enforced. In addition, one can check that, if $\mathbf{x}_{i,i+1} \in \mathbf{X}_i \cap \mathbf{X}_{i+1}$, then

$$\begin{aligned} d\mu_{i+1}^{\mathbf{X}_i \cap \mathbf{X}_{i+1}}(\mathbf{x}_{i,i+1}) &\stackrel{\text{def}}{=} \int_{\mathbf{y}_{i,i+1} \in \mathbf{X}_i^\perp \cap \mathbf{X}_{i+1}} d\mu_{i+1}(\mathbf{x}_{i,i+1} + \mathbf{y}_{i,i+1}) \\ &\stackrel{(15)}{=} \left(\int_{\mathbf{X}_i^\perp \cap \mathbf{X}_{i+1}} \mathbb{1}_{\mathbf{U}_{i+1}}(\mathbf{x}_{i,i+1} + \mathbf{y}_{i,i+1}) \right. \\ &\quad \left. \left(\int_{\mathbf{Z}_{i+1}} \prod_{j=i+2}^m \mathbb{1}_{\mathbf{U}_j} \circ \pi_{\mathbf{X}_j}(\mathbf{x}_{i,i+1} + \mathbf{y}_{i,i+1} + \mathbf{z}_{i+1}) d\mathbf{z}_{i+1} \right) d\mathbf{y}_{i,i+1} \right) d\mathbf{x}_{i,i+1} \\ &= \left(\int_{\mathbf{Z}_i} \prod_{j=i+1}^m \mathbb{1}_{\mathbf{U}_j} \circ \pi_{\mathbf{X}_j}(\mathbf{x}_{i,i+1} + \mathbf{z}_i) d\mathbf{z}_i \right) d\mathbf{x}_{i,i+1} \end{aligned}$$

since $(\mathbf{X}_i^\perp \cap \mathbf{X}_{i+1}) \oplus \mathbf{Z}_{i+1} = \mathbf{Z}_i$.

Thus, constraints (12) are satisfied. Moreover, they are saturated on \mathbf{U}_i . Eventually, one has

$$\mathbf{X} = \mathbf{X}_1 \oplus \mathbf{Z}_1$$

and thus

$$\begin{aligned} \int_{\mathbf{X}_1} d\mu_1 &= \int_{\mathbf{X}_1} \mathbb{1}_{\mathbf{U}_1}(\mathbf{x}_1) \left(\int_{\mathbf{Z}_1} \prod_{j=2}^m \mathbb{1}_{\mathbf{U}_j} \circ \pi_{\mathbf{X}_j}(\mathbf{x}_1 + \mathbf{z}_1) d\mathbf{z}_1 \right) d\mathbf{x}_1 \\ &= \int_{\mathbf{X}} \left(\prod_{j=1}^m \mathbb{1}_{\mathbf{U}_j} \circ \pi_{\mathbf{X}_j}(\mathbf{x}) \right) d\mathbf{x} \\ &= \int_{\mathbf{X}} \mathbb{1}_{\mathbf{K}}(\mathbf{x}) d\mathbf{x} \\ &= \text{vol } \mathbf{K}, \end{aligned}$$

that is, we have just proved that the value of the LP is larger than or equal to $\text{vol } \mathbf{K}$.

To prove the converse inequality, observe that our previous choice μ_1, \dots, μ_m saturates the constraints (12) while enforcing the constraints (14). Any other feasible solution $\tilde{\mu}_1, \dots, \tilde{\mu}_m$ directly satisfies the inequality $\tilde{\mu}_i \leq \mu_i$. In particular, $\tilde{\mu}_1 \leq \mu_1$ and thus

$$\int d\tilde{\mu}_1(\mathbf{x}_1) \leq \text{vol } \mathbf{K}.$$

□

Remark 9. The dual of the LP problem of Theorem 8 is the LP problem

$$\begin{aligned} \inf_{\substack{v_i \in \mathcal{L}_+(\mathbf{X}_i) \\ i=1, \dots, m}} & \int_{\mathbf{X}_m} v_m(\mathbf{x}_m) d\mathbf{x}_m \\ \text{s.t.} & v_1(\mathbf{x}_1) \geq 1 \quad \forall \mathbf{x}_1 \in \mathbf{U}_1 \\ & v_{i+1}(\mathbf{x}_{i+1}) \geq \int_{\mathbf{Y}_i} v_i(\mathbf{y}_i, \pi_{\mathbf{X}_i}(\mathbf{x}_{i+1})) d\mathbf{y}_i \quad \forall \mathbf{x}_{i+1} \in \mathbf{U}_{i+1}, i = 1, \dots, m-1. \end{aligned} \quad (16)$$

According to [11], there is no duality gap, i.e. the value of the dual is $\text{vol } \mathbf{K}$. For example, in the case of the bicylinder treated in Section 3.1, the dual reads:

$$\begin{aligned} \inf_{v_1, v_2 \in \mathcal{C}_+([0,1]^2)} \quad & \int_0^1 \int_0^1 v_2(x_2, x_3) dx_2 dx_3 \\ \text{s.t.} \quad & v_1(x_1, x_2) \geq 1 \quad \forall (x_1, x_2) \in \mathbf{U}_1 \\ & v_2(x_2, x_3) \geq \int_0^1 v_1(x_1, x_2) dx_1 \quad \forall (x_2, x_3) \in \mathbf{U}_2. \end{aligned}$$

Thus, if $(v_1^k, v_2^k)_{k \in \mathbb{N}}$ is a minimizing sequence for this dual LP, then the sets

$$\mathbf{A}^k := \left\{ (x_1, x_2, x_3) \in [0, 1]^3 : v_1^k(x_1, x_2) \geq 1, v_2^k(x_2, x_3) \geq \int_0^1 v_1^k(x_1, x_2) dx \right\}$$

are outer approximations of the set \mathbf{K} and the sequences $(\text{vol } \mathbf{A}^k)_k$ and $(\int v_2^k d\lambda^2)$ decrease to $\text{vol } \mathbf{K}$. Similar statements can be made for the general dual problem.

Corollary 10. [Convergence of the linear computation scheme]

The Moment-SOS hierarchy associated to problem (11) converges to $\text{vol } \mathbf{K}$.

Proof. We know that $\text{vol } \mathbf{K}$ is the value of the infinite dimensional primal LP on measures (11). The absence of duality gap (which is a consequence of the strong duality property in [4]) implies that there exists a sequence of continuous maps $(v_1^{(k)}, \dots, v_m^{(k)})_{k \in \mathbb{N}}$ feasible for its dual LP on functions (16) and such that $\int_{\mathbf{X}_m} v_m^{(k)}(\mathbf{x}_m) d\mathbf{x}_m \xrightarrow[k \rightarrow \infty]{} \text{vol } \mathbf{K}$. The Stone-Weierstrass Theorem allows to replace these continuous maps with vectors of polynomials while keeping the convergence of the optimization criterion. Eventually, Putinar's Positivstellensatz allows to replace the positivity constraints with SOS constraints while keeping the convergence of the optimization criterion. Thus, there is no relaxation gap between the dual (16) and its SOS reinforcement. Since the Moment-SOS hierarchy only consists of restricting the feasible set to bounded degree to enforce finite dimension, we can conclude that it converges to $\text{vol } \mathbf{K}$ when the upper bound on the degree tends to infinity. \square

Remark 11. The LP (11)-(14) is formulated as a single problem on m unknown measures. However, it is possible to split it in small chained subproblems to be solved in sequence. Each subproblem is associated with a maximal clique (in the linear clique tree) and it takes as input the results of the subproblem associated with its parent clique. This way, the sparse volume computation is split into m linked low-dimension problems and solved sequentially. This may prove useful when m is large because when solving the SDP relaxations associated with the single LP (11)-(14), the SDP solver may encounter difficulties in handling a high number of measures simultaneously. It should be easier to sequentially solve a high number of low-dimensional problems with only one unknown measure. Both formulations being strictly equivalent, this would not change the convergence properties of the sparse scheme.

3.3 Lower bounds for the volume

As explained in the introduction, the hierarchy of SDP relaxations associated with our infinite-dimensional LP provides us with a sequence of upper bounds on $\text{vol } \mathbf{K}$. One may also be interested in computing lower bounds on $\text{vol } \mathbf{K}$. In principle it suffices to apply the same methodology and approximate from above the volume of $\mathbf{B} \setminus \mathbf{K}$ since \mathbf{K} is included in the unit box \mathbf{B} . However, it is unclear whether $\mathbf{B} \setminus \mathbf{K}$ has also a sparse description. We show that this is actually the case and so one may exploit correlative sparsity to compute lower bounds although it is more technical. The following result is a consequence of Theorem 8:

Corollary 12. *If \mathbf{K} is sparse, then $\widehat{\mathbf{K}} := \mathbf{B} \setminus \mathbf{K}$ is sparse as well, and $\text{vol } \widehat{\mathbf{K}}$ is the value of the LP problem*

$$\begin{aligned} & \max_{\substack{\mu_{i,j} \in \mathcal{M}_+(\mathbf{X}_j) \\ 1 \leq i \leq j \leq m}} \sum_{j=1}^m \int_{\mathbf{X}_j} d\mu_{1,j} \\ \text{s.t.} \quad & \mu_{j,j} \leq \lambda^{m_j} \\ & \mu_{i,j} \leq \mu_{i+1,j}^{\mathbf{X}_i \cap \mathbf{X}_{i+1}} \otimes \lambda^{n_i} \quad i = 1, \dots, j-1 \\ & \text{spt } \mu_{i,j} \subset \mathbf{U}_i \quad i = 1, \dots, j-1 \\ & \text{spt } \mu_{j,j} \subset \text{cl } \widehat{\mathbf{U}}_j \end{aligned}$$

where $m_j := \dim \mathbf{X}_j$, $n_i := \dim \mathbf{X}_{i+1}^\perp \cap \mathbf{X}_i$, $\widehat{\mathbf{U}}_j := [0, 1]^{m_j} \setminus \mathbf{U}_j$ is open and $\text{cl } \widehat{\mathbf{U}}_j$ denotes its closure⁵.

Proof. The following description

$$\widehat{\mathbf{K}} = \bigsqcup_{j=1}^m \left[\bigcap_{i=1}^{j-1} \mathbf{U}_i \cap \widehat{\mathbf{U}}_j \right],$$

where \bigsqcup stands for disjoint union, is sparse. Indeed the description of the basic semi-algebraic set

$$\mathbf{V}_j := \bigcap_{i=1}^{j-1} \mathbf{U}_i \cap \widehat{\mathbf{U}}_j$$

is sparse. In addition, by σ -additivity of the Lebesgue measure, it holds

$$\text{vol } \widehat{\mathbf{K}} = \sum_{j=1}^m \text{vol } \mathbf{V}_j.$$

Finally, by using Theorem 8 we conclude that $\text{vol } \mathbf{V}_j$ is the value of LP consisting of maximizing $\int_{\mathbf{X}_j} d\mu_{1,j}$ subject to the same constraints as in the above LP problem. Summing up yields the correct value. \square

4 Accelerating convergence

As already mentioned, the convergence of the standard SDP relaxations (3) for solving the GMP (1) is expected to be slow in general. To cope with this issue we introduce additional linear constraints that are redundant for the infinite dimensional GMP, and that are helpful to accelerate the convergence of the SDP relaxations. These constraints come from a specific application of Stokes' theorem.

4.1 Dense Stokes constraints

We first focus on the dense formulation (1). We know that the optimal measure of our infinite dimensional LP is $\mu = \lambda_{\mathbf{K}}$. Thus, one can put additional constraints in the hierarchy in order to give more information on the target sequence of moments, without increasing the dimension of the SDP relaxation. To keep the optimal value unchanged, such constraints should be redundant in the infinite dimensional LP. Ideally, we would like to characterize the whole set of polynomials p such that

$$\int_{\mathbf{K}} p(\mathbf{x}) d\mu(\mathbf{x}) = 0.$$

⁵This is necessary since the support of a measure is a closed set by definition.

Indeed, given any such polynomial p , the moments \mathbf{m} of μ necessarily satisfy the linear constraint $L_{\mathbf{m}}(p) = 0$. However, for a general semi-algebraic set \mathbf{K} , we do not have an explicit description of this set of polynomials. Nevertheless, we can generate many of them, and hence improve convergence of the SDP relaxations significantly, as it was done originally in [20] in another context. Let us explain how we generate these linear moment constraints.

We recall that Stokes' theorem states that if \mathbf{O} is an open set of \mathbb{R}^n with a boundary $\partial\mathbf{O}$ smooth almost everywhere, and ω is a $(n-1)$ -differential form on \mathbb{R}^n , then one has

$$\int_{\partial\mathbf{O}} \omega = \int_{\mathbf{O}} d\omega.$$

A corollary to this theorem is obtained by choosing $\omega(\mathbf{x}) = \mathbf{h}(\mathbf{x}) \cdot \mathbf{n}(\mathbf{x}) \sigma(d\mathbf{x})$, where the dot denotes the inner product between \mathbf{h} , a smooth (ideally polynomial) vector field, and \mathbf{n} the outward pointing unit vector orthogonal to the boundary $\partial\mathbf{O}$, and σ denotes the $(n-1)$ -dimensional Hausdorff measure on $\partial\mathbf{O}$. In this case, one obtains the Gauss formula [21]

$$\int_{\partial\mathbf{O}} \mathbf{h}(\mathbf{x}) \cdot \mathbf{n}(\mathbf{x}) \sigma(d\mathbf{x}) = \int_{\mathbf{O}} \operatorname{div} \mathbf{h}(\mathbf{x}) d\mathbf{x}.$$

Choosing $\mathbf{h}(\mathbf{x}) = h(\mathbf{x}) \mathbf{e}_i$, where h is a smooth function (ideally a polynomial) and \mathbf{e}_i is the i -th vector of the canonical basis of \mathbb{R}^n , one obtains the following vector equality

$$\int_{\partial\mathbf{O}} h(\mathbf{x}) \mathbf{n}(\mathbf{x}) d\sigma(\mathbf{x}) = \int_{\mathbf{O}} \mathbf{grad} h(\mathbf{x}) d\mathbf{x}.$$

Then, if we choose $\mathbf{O} = \mathbf{K} \setminus \partial\mathbf{K}$ and a polynomial h vanishing on $\partial\mathbf{K}$, the vector constraint

$$\int_{\mathbf{K}} \mathbf{grad} h(\mathbf{x}) d\mathbf{x} = 0$$

is automatically satisfied and it can be added to the optimization problem (1) without changing its optimal value. Such constraints are redundant in the infinite-dimensional LP formulation (1) but not in the SDP relaxations (3). It has been numerically shown in [20] that adding these constraints dramatically increases the convergence rate of the hierarchy of SDP relaxations.

These so-called *Stokes constraints* can be added to the formulation of problem (1) to yield

$$\begin{aligned} \max_{\mu, \hat{\mu} \in \mathcal{M}_+(\mathbf{B})} \quad & \int d\mu \\ \text{s.t.} \quad & \mu + \hat{\mu} = \lambda_{\mathbf{B}} \\ & (\mathbf{grad} h)\mu = \mathbf{grad}(h\mu) \\ & \operatorname{spt} \mu \subset \mathbf{K} \\ & \operatorname{spt} \hat{\mu} \subset \mathbf{B} \end{aligned} \tag{17}$$

where h is any polynomial vanishing on the boundary of \mathbf{K} , without changing its value $\operatorname{vol} \mathbf{K}$. The vector constraint $(\mathbf{grad} h)\mu = \mathbf{grad}(h\mu)$ should be understood in the sense of distributions, i.e. for all test functions $v \in \mathcal{C}^1(\mathbf{B})$ it holds

$$\int (\mathbf{grad} h)v d\mu = - \int (\mathbf{grad} v)h d\mu$$

or equivalently

$$\int ((\mathbf{grad} h)v + (\mathbf{grad} v)h) d\mu = \int \mathbf{grad}(hv) d\mu = 0$$

which becomes a linear moment constraint

$$L_{\mathbf{m}}(\mathbf{grad}(hv)) = 0$$

if v is polynomial. In practice, when implementing the SDP relaxation of degree d , we choose $h(\mathbf{x}) := \prod_{j=1}^m g_j(\mathbf{x})$ and $v(\mathbf{x}) = \mathbf{x}^\alpha$, $\alpha \in \mathbb{N}^n$, $|\alpha| \leq d+1 - \sum_{j=1}^m \deg g_j$.

Remark 13. The dual to the LP problem (17) is the LP problem

$$\begin{aligned} & \inf_{\substack{v \in \mathcal{C}_+(\mathbf{B}) \\ \mathbf{w} \in \mathcal{C}(\mathbf{B})^n}} \int v \, d\lambda_{\mathbf{B}} \\ & \text{s.t. } v + \operatorname{div}(h\mathbf{w}) \geq \mathbb{1}_{\mathbf{K}}. \end{aligned}$$

It follows that the function v is not required anymore to approximate from above the discontinuous indicator function $\mathbb{1}_{\mathbf{K}}$, so that the Gibbs effect is reduced. We believe that the infimum in this dual of Stokes is in fact a minimum, i.e. there exists optimal decision variables v, \mathbf{w} . This would make the Gibbs effect totally disappear. Proving this reduces to a problem of existence and uniqueness of the solution to a degenerate linear PDE, and it is out of the scope of this paper.

4.2 Sparse Stokes constraints: the bicylinder

We have designed efficient Stokes constraints for the dense formulation of problem (1), at the price of introducing in problem (4) a polynomial h vanishing on the boundary of \mathbf{K} . However, in the sparse case (7), the polynomial h would destroy the sparsity structure, as it is the product of all polynomials defining \mathbf{K} . So we must adapt our strategy to introduce sparse Stokes constraints.

In this section, to keep the notations simple, we illustrate the ideas on our introductory bicylinder example of Section 3.1. Considering the optimal measures μ_1 and μ_2 defined in (8),(9), we can apply Stokes constraints derived from the Gauss formula, in the directions in which they are Lebesgue: for μ_1 in the x_1 direction and for μ_2 in the remaining directions. To see this, define

$$\begin{aligned} \mathbf{h}_1(x_1, x_2) &= g_1(x_1, x_2) \mathbf{e}_1, \\ \mathbf{h}_2(x_2, x_3) &= g_2(x_2, x_3) \mathbf{e}_2, \\ \mathbf{h}_3(x_2, x_3) &= g_2(x_2, x_3) \mathbf{e}_3 \end{aligned}$$

where $g_i(x_i, x_{i+1}) = 1 - x_i^2 - x_{i+1}^2$, such that $\mathbf{h}_1 \cdot \mathbf{n}_{\mathbf{U}_1}$ vanishes on the boundary of \mathbf{U}_1 and $\mathbf{h}_2 \cdot \mathbf{n}_{\mathbf{U}_2}$ and $\mathbf{h}_3 \cdot \mathbf{n}_{\mathbf{U}_2}$ vanish on the boundary of \mathbf{U}_2 , where $\mathbf{n}_{\mathbf{U}_i}$ is the outward point vector orthogonal to the boundary of \mathbf{U}_i . For $i, j, k \in \mathbb{N}$, the Gauss formula yields

$$\begin{aligned} \int_{\mathbf{U}_1} \frac{\partial}{\partial x_1} (g_1(x_1, x_2) x_1^i x_2^j) \, d\mu_1 &= 0, \\ \int_{\mathbf{U}_2} \frac{\partial}{\partial x_2} (g_2(x_2, x_3) x_2^j x_3^k) \, d\mu_2 &= 0, \\ \int_{\mathbf{U}_2} \frac{\partial}{\partial x_3} (g_2(x_2, x_3) x_2^j x_3^k) \, d\mu_2 &= 0. \end{aligned}$$

Hence, adding these constraints does not change the optimal value of the LP problem (7).

4.3 General sparse Stokes constraints

Consider the sequential decomposition of Theorem 8:

$$\begin{aligned} & \max_{\mu_i \in \mathcal{M}_+(\mathbf{X}_i)} \int d\mu_i \\ & \text{s.t. } \mu_i \leq \mu_{i+1}^{\mathbf{X}_i \cap \mathbf{X}_{i+1}} \otimes \lambda^{n_i} \\ & \text{spt } \mu_i \subset \mathbf{U}_i \end{aligned}$$

for $1 \leq i \leq m - 1$, and

$$\begin{aligned} \max_{\mu_m \in \mathcal{M}_+(\mathbf{X}_m)} \quad & \int d\mu_m \\ \text{s.t.} \quad & \mu_m \leq \lambda^{n_m} \\ & \text{spt } \mu_m \subset \mathbf{U}_m. \end{aligned}$$

Our algorithm consists of sequentially solving these problems, starting with determining μ_m , then μ_{m-1} , and so on until μ_1 , whose mass will be $\text{vol}(\mathbf{K})$. We implement Stokes constraints on each one of these problems. For the problem in μ_m , we implement regular Stokes constraints as in section 4.1:

$$\begin{aligned} \max_{\mu_m \in \mathcal{M}_+(\mathbf{X}_m)} \quad & \int d\mu_m \\ \text{s.t.} \quad & \mu_m \leq \lambda^{n_m} \\ & (\mathbf{grad} h_m)\mu_m = (\mathbf{grad} h_m \mu_m) \\ & \text{spt } \mu_m \subset \mathbf{U}_m \end{aligned}$$

where h_m is a polynomial vanishing on $\partial\mathbf{U}_m$.

Then, let $i \in \{1, \dots, m - 1\}$ and suppose that μ_{i+1} is known, such that determining μ_i is reduced to a linear programming problem. From the arguments of Section 3, we know that the optimal measure μ_i is supported on \mathbf{U}_i and that on this set it is the product measure between $\mu_{i+1}^{\mathbf{X}_i \cap \mathbf{X}_{i+1}}$ and the uniform measure on $\mathbf{Y}_i = \langle x_j \rangle_{j \in C_i \cap C_{i+1}^c}$. Since Stokes' theorem is only valid for uniform measures, it will only apply to $\mu_i^{\mathbf{Y}_i}$.

Let $J \subset \{1, \dots, n\}$. For $f \in \mathcal{C}^1(\mathbb{R}^n)$ we define

$$\nabla_J f := \left(\frac{\partial f}{\partial x_j} \right)_{j \in J}$$

such that $\nabla_{\{1, \dots, n\}} f = \mathbf{grad} f$ and $\nabla_{\{j\}} f = \frac{\partial f}{\partial x_j}$ for example. This notation allows us to define Stokes constraints exactly in the directions we are interested in and to formulate the general sparse Stokes constraints:

$$\begin{aligned} \max_{\mu_i \in \mathcal{M}_+(\mathbf{X}_i)} \quad & \int d\mu_i \\ \text{s.t.} \quad & \mu_i \leq \mu_{i+1}^{\mathbf{X}_i \cap \mathbf{X}_{i+1}} \otimes \lambda^{n_i} \\ & (\nabla_{C_i \cap C_{i+1}^c} h_i)\mu_i = (\nabla_{C_i \cap C_{i+1}^c} h_i \mu_i) \\ & \text{spt } \mu_i \subset \mathbf{U}_i \end{aligned}$$

where h_i is a polynomial vanishing on $\partial\mathbf{U}_i$.

Remark 14. In some cases, in both dense and sparse contexts, these Stokes constraints can be slightly improved by choosing a different polynomial h_j for each basis vector \mathbf{e}_j when applying the Gauss formula, such that h_j can be taken with the lowest possible degree, allowing for a better implementation of the hierarchy. For example, if one is looking for the volume of $\mathbf{K} := [0, 1]^2$, the polynomial vanishing on $\partial\mathbf{K}$ with the lowest degree is $h(x_1, x_2) := x_1(1 - x_1)x_2(1 - x_2)$, but one can formulate Stokes constraints by applying the Gauss formula to $x_1(1 - x_1)\mathbf{e}_1$ and $x_2(1 - x_2)\mathbf{e}_2$, instead of $h(x_1, x_2)\mathbf{e}_1$ and $h(x_1, x_2)\mathbf{e}_2$. By doing so, one would replace the constraint $(\mathbf{grad} h)\mu = \mathbf{grad}(h\mu)$ with $\left(\frac{\partial h_j}{\partial x_j}\right)\mu = \frac{\partial}{\partial x_j}(h_j\mu)$ for every possible j . This is what we actually implemented in our numerical examples, but we presented the Stokes constraints in the restrictive case of $h_j = h$ for all j for the sake of readability.

5 Numerical examples

5.1 Bicylinder revisited

We refer to (1) as the dense problem and to (7) as the sparse problem. For both problems, we consider instances with and without additional Stokes constraints. Note that for the bicylinder example of Section 3.1 the optimal value for both the dense and the sparse problem is

$$\text{vol } \mathbf{K} = \frac{16}{3} \approx 5.3333$$

since adding Stokes constraints does not change the optimal value.

We solve the SDP relaxations with Mosek on a standard laptop, for various relaxation orders and we report the bounds and the computation times in Table 1. We observe a slow convergence

d	full		sparse	
	without Stokes	with Stokes	without Stokes	with Stokes
2	7.8232 (1.0s)	5,828 (1.1s)	7,7424 (1.1s)	5,4984 (1.1s)
3	7.2368 (0.9s)	5,4200 (1.3s)	7,1920 (0.9s)	5,3488 (1.1s)
4	7.0496 (1.4s)	5,3520 (2.2s)	7,0040 (1.2s)	5,3376 (1.2s)
5	6,8136 (3.1s)	5,3400 (4.4s)	6,7944 (1.8s)	5,3352 (1.8s)
6	6,7376 (7.2s)	5,3376 (8.2s)	6,6960 (2.1s)	5,3344 (2.3s)
7	6,6336 (12.8s)	5,3360 (18.3s)	6,6168 (2.6s)	5,3344 (3.2s)

Table 1: Bounds on the volume (and computation times in seconds) vs relaxation order for the bicylinder.

for the dense and the sparse versions without Stokes constraints, and a much faster convergence with Stokes constraints. We also observe significantly smaller computation times when using the sparse formulation.

5.2 A nonconvex set

Let $\mathbf{X} := \mathbb{R}^5$, $\mathbf{X}_1 = \langle x_1, x_2 \rangle$, $\mathbf{X}_2 = \langle x_1, x_3 \rangle$, $\mathbf{X}_3 = \langle x_1, x_4 \rangle$, $\mathbf{X}_4 = \langle x_1, x_5 \rangle$ and

- $\mathbf{g}_i(x_1, x_{i+1}) := (2x_1^2 - x_{i+1}^2 - 1, x_1(1 - x_1), x_{i+1}(1 - x_{i+1}))$, $i = 1, \dots, 4$
- $\mathbf{U}_i := \mathbf{g}_i^{-1}((\mathbb{R}_+)^3) = \{(x_1, x_{i+1}) \in [0, 1]^2 : 2x_1^2 - x_{i+1}^2 \geq 1\}$, $i = 1, \dots, 4$.

Let us approximate the volume of the sparse set

$$\mathbf{K} := \{(x_1, x_2, x_3, x_4, x_5) \in [0, 1]^5 : 2x_1^2 - x_{i+1}^2 \geq 1, i = 1, \dots, 4\} = \bigcap_{i=1}^4 \pi_{\mathbf{X}_i}^{-1}(\mathbf{U}_i).$$

Here the coordinates x_2, x_3, x_4 and x_5 do not interact: they are only linked with the coordinate x_1 . The proper way to apply our linear computation Theorem 8 is to define a linear clique tree as shown in Figure 6.




Figure 6: Graph with linear clique tree for the nonconvex set.

This yields the following formulation

$$\begin{aligned}
\text{vol } \mathbf{K} &= \max_{\substack{\mu_i \in \mathcal{M}_+(\mathbf{X}_i) \\ i=1, \dots, 4}} \int_{\mathbf{X}_1} d\mu_1 & (18) \\
\text{s.t.} \quad & d\mu_1(x_1, x_2) \leq d\mu_2^{(x_1)}(x_1) dx_2 \\
& d\mu_2(x_1, x_3) \leq d\mu_3^{(x_1)}(x_1) dx_3 \\
& d\mu_3(x_1, x_4) \leq d\mu_4^{(x_1)}(x_1) dx_4 \\
& d\mu_4(x_1, x_5) \leq dx_1 dx_5 \\
& \text{spt } \mu_i \subset \mathbf{U}_i \quad i = 1, \dots, 4
\end{aligned}$$

with Stokes constraints

$$\begin{aligned}
\frac{\partial}{\partial x_2} \left[(2x_1^2 - x_2^2 - 1) x_2 (1 - x_2) \right] d\mu_1(x_1, x_2) &= \frac{\partial}{\partial x_2} \left[(2x_1^2 - x_2^2 - 1) x_2 (1 - x_2) d\mu_1(x_1, x_2) \right] \\
\frac{\partial}{\partial x_3} \left[(2x_1^2 - x_3^2 - 1) x_3 (1 - x_3) \right] d\mu_2(x_1, x_3) &= \frac{\partial}{\partial x_3} \left[(2x_1^2 - x_3^2 - 1) x_3 (1 - x_3) d\mu_2(x_1, x_3) \right] \\
\frac{\partial}{\partial x_4} \left[(2x_1^2 - x_4^2 - 1) x_4 (1 - x_4) \right] d\mu_3(x_1, x_4) &= \frac{\partial}{\partial x_4} \left[(2x_1^2 - x_4^2 - 1) x_4 (1 - x_4) d\mu_3(x_1, x_4) \right] \\
\frac{\partial}{\partial x_5} \left[(2x_1^2 - x_5^2 - 1) x_5 (1 - x_5) \right] d\mu_4(x_1, x_5) &= \frac{\partial}{\partial x_5} \left[(2x_1^2 - x_5^2 - 1) x_5 (1 - x_5) d\mu_4(x_1, x_5) \right] \\
\frac{\partial}{\partial x_1} \left[(2x_1^2 - x_5^2 - 1) x_1 (1 - x_1) \right] d\mu_4(x_1, x_5) &= \frac{\partial}{\partial x_1} \left[(2x_1^2 - x_5^2 - 1) x_1 (1 - x_1) d\mu_4(x_1, x_5) \right].
\end{aligned}$$

We can compute analytically

$$\text{vol } \mathbf{K} = \frac{1}{15} (7 - 4\sqrt{2}) \simeq 0.0895.$$

On Figure 7 we show results from solving several relaxations via the dense and the sparse approach, with and without Stokes constraints. While solving with Mosek the degree 12 dense relaxation took about 1000 seconds, solving the degree 12 sparse relaxation took less than 10 seconds. With the sparse relaxations it was possible to go much higher in the hierarchy. Figure 7b shows convincingly how Stokes constraints accelerate the convergence of the hierarchy. We can also observe that the nonconvexity of \mathbf{K} poses no special difficulty for the volume computation.

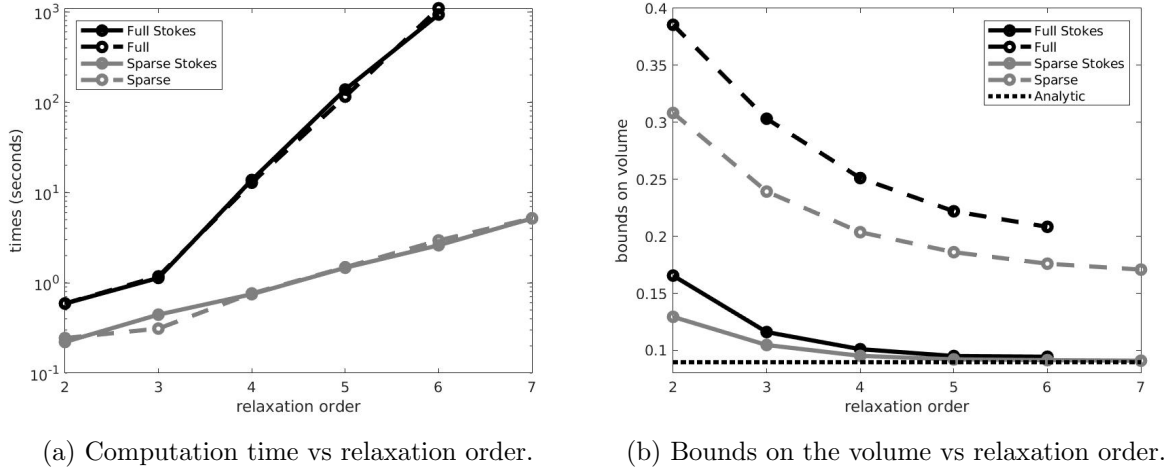


Figure 7: Performance for the nonconvex set.

5.3 A high dimensional polytope

Consider

$$\mathbf{K}_n := \{\mathbf{x} \in [0, 1]^n : x_i + x_{i+1} \leq 1, i = 1, \dots, n-1\}.$$

According to [22], for any $\theta \in]-\frac{\pi}{2}, \frac{\pi}{2}[$, one has the elegant formula :

$$\tan \theta + \sec \theta = 1 + \sum_{n=1}^{\infty} \text{vol } \mathbf{K}_n \theta^n$$

which allows to compute analytically the volume for n arbitrarily large. For example when $n = 20$ we obtain

$$\text{vol } \mathbf{K}_{20} = \frac{14814847529501}{97316080327065600} \approx 1.522 \cdot 10^{-4}.$$

From the SDP viewpoint, $\text{vol } \mathbf{K}_n$ is computed by solving relaxations of the LP problem given in Theorem 8 where $m = n-1$, $\mathbf{X}_i = \langle x_i, x_{i+1} \rangle$ and $\mathbf{g}_i(x_i, x_{i+1}) = (x_i, x_{i+1}, 1 - x_i - x_{i+1})$, $i = 1, \dots, n-1$.

We implemented the volume computation algorithm for $n = 20$, with Stokes constraints. This cannot be achieved without resorting to sparse computation as the dimension is too high for regular SDP solvers. With the sparse formulation however we could solve relaxations up to degree 28 in less than 100 seconds, see Figure 8. Note however, that the analytic volume is of the order of 10^{-4} . In consequence we observe a non monotonicity of the relaxation values which contradicts the theory. This issue is surprising as the Mosek SDP solver terminates without reporting issues. This indicates that computing small volumes in large dimension can be numerically sensitive.

In order to fix the monotonicity issue, we added a sparse rescaling to our problem. The idea is the following: at each step of the algorithm, the mass of the measure μ_i is less than the mass of the reference measure

$$\rho_i := \mu_{i+1}^{\mathbf{X}_i \cap \mathbf{X}_{i+1}} \otimes \lambda^{n_i}.$$

Defining

$$\epsilon_i := \frac{|\mu_i|}{|\rho_i|} \in]0, 1[,$$

we obtain that

$$\text{vol } \mathbf{K} = \prod_{i=1}^m \epsilon_i$$


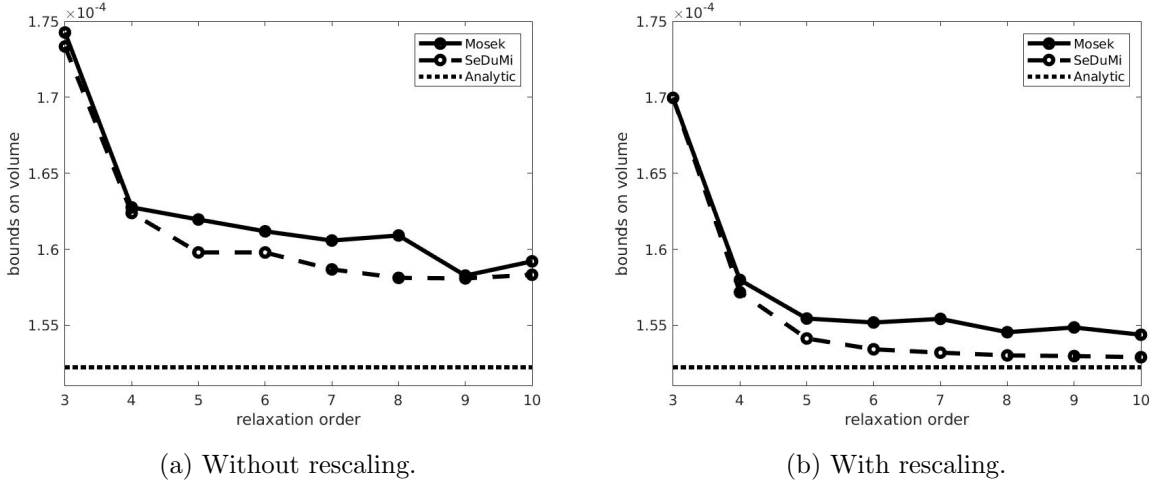


Figure 8: Bounds on the volume vs relaxation order for the high dimensional polytope.



(a) Without rescaling.

(b) With rescaling.

Figure 9: Bounds on the volume vs relaxation order for the high dimensional polytope.

as a telescoping product, since $|\rho_m| = \text{vol } \mathbf{B} = 1$. As a result, if m is large and the ϵ_i are small, one can expect the volume to be very small, which explains why the SDP solver encounters difficulties. Thus, a solution is to multiply each domination constraint by a well-chosen rescaling factor ϵ such that the mass of μ_i does not decrease too much. The resulting LP is as follows

$$\begin{aligned}
 \text{vol } \mathbf{K} &= \epsilon^{m-1} \max_{\substack{\mu_i \in \mathcal{M}_+(\mathbf{X}_i) \\ i=1, \dots, m}} \int_{\mathbf{X}_1} d\mu_1 & (19) \\
 \text{s.t. } & \epsilon \mu_i \leq \mu_{i+1}^{\mathbf{X}_i \cap \mathbf{X}_{i+1}} \otimes \lambda^{n_i} & i = 1, \dots, m-1 \\
 & \mu_m \leq \lambda^{n_m} \\
 & \text{spt } \mu_i \subset \mathbf{U}_i & i = 1, \dots, m.
 \end{aligned}$$

Figure 9 gives a comparison between the results obtained with and without sparse rescaling, using the SDP Solvers SeDuMi and Mosek, for the choice $\epsilon = \frac{1}{2}$.

First, one can see that without rescaling (Figure 9a), both SeDuMi and Mosek have accuracy issues that make them lose monotonicity, while the rescaling (Figure 9b) allows to recover monotonicity at least when using SeDuMi (which is slower but more accurate than Mosek to our general experience). Second, it is clear that the relative approximation error is much smaller with scaling. This, combined to the fact that the error is relative (after rescaling, the error is

much smaller), demonstrates the power of our rescaling method.

5.4 A nonconvex high dimensional set

Finally, we consider the set already mentioned in the introduction, which is both nonconvex and high dimensional. Let

$$\mathbf{K}_n := \{\mathbf{x} \in [0, 1]^n : x_{i+1} x_i \leq 1/2, i = 1, 2, \dots, n-1\}$$

whose analytic volume is a function of the dimension n . For $n = 3$ the analytic volume is 0.75, for $n = 4$ it is 0.6566, approximately. In higher dimensions we do not have an analytic expression for the volume. However, in order to get a feeling for its value for bigger n , we ran a Monte Carlo simulation⁶ with one million samples for $n = 10, 20, 50$, and 100.

Before we go on, let us emphasize that the method proposed in this paper is not in concurrence with the Monte Carlo approach. While the Monte Carlo gives a *probabilistic* estimate of the volume, our method provides a *guaranteed* upper bound. Nonetheless, it would be concerning if the computed upper bound was much smaller than the confidence interval, and we consider our approximation valid, when it returns something in the order of the Monte Carlo approximation. The results for different dimensions n and solved with the Mosek SDP solver are summarized in Table 2. As in the previous section we experience accuracy issues for the relaxations of order 14 and $n = 20, 100$, as well as for order 16 and $n = 50$. Otherwise, the approximations provide better upper bounds for increased relaxation orders as expected. For $n = 3, 4$ the approximation is reasonably close to the analytic value. For $n = 10, 20, 50$ our scheme provides an upper bound in the same order of magnitude as the 99%-confidence interval of the Monte Carlo simulation. We interpret this as a validation for both the Monte Carlo approach and our own approach. For $n = 100$ we could not derive a meaningful confidence interval. Indeed, as our approximation shows, the volume for $n = 100$ is less than $9 \cdot 10^{-6}$. In order to get an accuracy of $\varepsilon = 10^{-6}$ one would have to draw approximately $N = \frac{1}{\varepsilon^2} = 10^{12}$ samples. With our non-sophisticated implementation, the Monte Carlo simulation for one million points took about 5 seconds. Extending this linearly to a simulation with 10^{12} samples would therefore take a little less than 2 months ($5 \cdot 10^6 s \simeq 1389 h \simeq 58 d$). With more sophisticated methods, this time could certainly be reduced dramatically. However, it sets the 44 minutes it took to solve relaxation order 16 for $n = 100$ into perspective.

6 General sparse volume computation

6.1 General correlative sparsity pattern

Let us describe a general method to compute the volume of

$$\mathbf{K} := \bigcap_{i=1}^m \mathbf{K}_i$$

where $\mathbf{K}_i = \{\mathbf{x} \in \mathbf{X} : \mathbf{g}_i(\mathbf{x}) \geq 0\}$ and $(\mathbf{g}_1, \dots, \mathbf{g}_m)$ is a correlative sparse family of polynomial vectors with associated coordinate subspace decomposition $\mathbf{X} = \sum_{i=1}^m \mathbf{X}_i$. For this we construct the correlation graph $G = (V, E)$ as follows:

- $V = \{1, \dots, n\}$ represents the canonical basis $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ of \mathbf{X} ;
- $E = \{(i, j) \in \{1, \dots, n\}^2 : i \neq j \ \& \ \mathbf{e}_i, \mathbf{e}_j \in \mathbf{X}_k \text{ for some } k \in \{1, \dots, m\}\}$.

⁶We provide a quick introduction in Appendix B.

d	n=3		n = 4		n=10	
	value	time (s)	value	time (s)	value	time (s)
4	7.86E-01	0.95	7.09E-01	0.61	3.93E-01	1.24
5	7.73E-01	2.87	6.90E-01	0.78	3.57E-01	1.84
6	7.69E-01	2.74	6.84E-01	0.86	3.45E-01	4.21
7	7.66E-01	4.58	6.79E-01	1.83	3.38E-01	4.55
8	7.63E-01	5.00	6.77E-01	2.29	3.34E-01	5.97
9	7.63E-01	6.11	6.74E-01	3.33	3.30E-01	11.56
10	7.62E-01	9.83	6.73E-01	6.86	3.26E-01	18.21
11	7.61E-01	18.16	6.72E-01	8.57	3.26E-01	22.24
12	7.60E-01	19.45	6.71E-01	10.43	3.23E-01	33.78
13	7.60E-01	22.49	6.70E-01	17.89	3.22E-01	74.00
14	7.60E-01	27.02	6.69E-01	26.84	3.21E-01	79.68
15	7.59E-01	32.90	6.69E-01	39.25	3.20E-01	119.7
16	7.58E-01	78.20	6.68E-01	61.32	3.19E-01	176.6
ana/mc	7.50E-01	-	6.57E-01	-	[2.99e-01, 3.03e-01]	
d	n = 20		n=50		n = 100	
	value	time (s)	value	time (s)	value	time (s)
4	1.47E-01	5.11	7.68E-03	10.64	9.49E-05	15.19
5	1.20E-01	3.58	4.78E-03	15.10	4.80E-05	26.87
6	1.11E-01	8.13	3.86E-03	21.75	2.84E-05	49.89
7	1.07E-01	11.06	3.42E-03	48.31	2.27E-05	77.07
8	1.03E-01	17.93	3.20E-03	72.78	1.91E-05	135.08
9	1.00E-01	33.31	2.99E-03	120.49	1.63E-05	202.12
10	9.81E-02	41.02	2.99E-03	103.61	1.44E-05	299.44
11	9.70E-02	85.83	2.89E-03	165.56	1.22E-05	441.67
12	9.59E-02	117.08	2.77E-03	220.84	1.19E-05	623.24
13	9.51E-02	138.38	2.67E-03	314.00	1.08E-05	850.92
14	9.57E-02	156.32	2.60E-03	457.92	1.10E-05	1175.02
15	9.39E-02	249.82	2.54E-03	685.64	9.86E-06	1589.49
16	9.36E-02	357.87	2.56E-03	859.60	9.46E-06	2623.02
ana/mc	[8.09e-02, 8.24e-02]		[1.48e-03, 1.68e-03]		-	

Table 2: **A nonconvex high dimensional set:** ana/mc refers to the analytic value and the 99%-confidence interval, respectively.

As stated in Section 2, we suppose that the correlation graph of $(\mathbf{g}_i)_{1 \leq i \leq m}$ has exactly m maximal cliques (see Section 6.5 for discussions when it is not the case) that are in correspondence with the \mathbf{X}_i .

Let \mathcal{K} be the set of maximal cliques of G . We will use the following property of graphs:

Definition 15. (CIP) The graph $G = (V, E)$ is said to satisfy the *clique intersection property (CIP)* iff there is a clique tree $T = (\mathcal{K}, \mathcal{E})$, such that for all $C, C' \in \mathcal{K}$, $C \cap C' \subset C''$ for any C'' on the path connecting C and C' in the tree T .

Such a property ensures that G is chordal⁷, see [19]. We then replace Assumption 7 with the following strong correlative sparsity assumption:

Assumption 16. (DIP) We suppose that there is a clique tree $T = (\mathcal{K}, \mathcal{E})$, rooted in some C_1 , that simultaneously satisfies the CIP and the following disjoint intersection property (DIP): $\forall C, C', C'' \in \mathcal{K}$, if $(C, C') \in \mathcal{E}_T$ and $(C, C'') \in \mathcal{E}$ then $C' = C''$ or $C' \cap C'' = \emptyset$.

In words, each clique has an empty intersection with all its siblings. See Appendix A for details on how to check this assumption and construct such a tree when it exists. See Section 6.5 for possible solutions when Assumption 16 does not hold. Figure 10 illustrates the meaning of this assumption for $n = 12$ and $m = 8$. One can check that Assumptions 6 and 16 hold.

Remark 17. With these assumptions, the only possible clique trees for applying our method to the nonconvex example illustrated in Figure 6 are linear clique trees. Indeed, any branched clique tree would imply sibling cliques containing x_1 .




Figure 10: Chordal graph (left) with its clique tree (right).

⁷The CIP always holds, up to a chordal extension. In particular, cyclic graphs can be handled with empty interactions between well-chosen variables.

6.2 Distributed computation theorem

One can formulate a simple generalization of the sequential implementation of Theorem 8 to our general correlative sparsity pattern.

Theorem 18. *Let Assumptions 6 and 16 hold. Let $T = (\mathcal{K}, \mathcal{E})$ be a clique tree as in Assumption 16. Then $\text{vol } \mathbf{K} = \int_{\mathbf{X}_1} d\mu_1^*$ where for $i = 1, \dots, m$, μ_i^* is an optimal solution to*

$$\max_{\mu_i \in \mathcal{M}_+(\mathbf{X}_i)} \int_{\mathbf{X}_i} d\mu_i \quad (20)$$

$$\text{s.t.} \quad \mu_i \leq \left(\bigotimes_{(C_i, C_j) \in \mathcal{E}} \mu_j^{*\mathbf{X}_i \cap \mathbf{X}_j} \right) \otimes \lambda^{n_i} \quad (21)$$

$$\text{spt } \mu_i \subset \mathbf{U}_i \quad (22)$$

$$\text{and } n_i = \dim \mathbf{X}_i \cap \left(\sum_{(C_i, C_j) \in \mathcal{E}} \mathbf{X}_j \right)^\perp = \left| C_i \cap \left(\bigcup_{(C_i, C_j) \in \mathcal{E}} C_j \right)^c \right|.$$

Proof. We define, for $i = 1, \dots, m$,

$$\mathbf{Y}_i := \mathbf{X}_i \cap \left(\sum_{(C_i, C_j) \in \mathcal{E}} \mathbf{X}_j \right)^\perp = \langle x_k \rangle_{k \in C_i; (C_i, C_j) \in \mathcal{E} \Rightarrow k \notin C_j}$$

and we observe that, according to Assumption 16, for any $i = 1, \dots, m$

$$\mathbf{X}_i = \left(\bigoplus_{(C_i, C_j) \in \mathcal{E}} \mathbf{X}_i \cap \mathbf{X}_j \right) \oplus \mathbf{Y}_i.$$

Thus, constraint (21) is well-posed.

For $i = 1, \dots, m$, let $\mathcal{D}(i) := \{j \neq i : \exists \text{ an oriented path from } C_i \text{ to } C_j \text{ in } T\}$, the set of descendants of C_i , as well as $\mathbf{Z}_i := \sum_{j \in \mathcal{D}(i)} \mathbf{X}_j = \langle x_k \rangle_{k \in C_j; j \in \mathcal{D}(i)}$ and $q_i := \dim \mathbf{Z}_i = \left| \bigcup_{j \in \mathcal{D}(i)} C_j \right|$.

We are going to show by induction that for $i = 1, \dots, m$,

$$\mu_i^* = \mathbb{1}_{\mathbf{U}_i} \left(\prod_{j \in \mathcal{D}(i)} \mathbb{1}_{\mathbf{U}_j} \circ \pi_{\mathbf{X}_j} \lambda^{q_j} \right)^{\mathbf{X}_i \cap \mathbf{Z}_i} \otimes \lambda^{n_i}.$$

Our base cases are the leaves of T , i.e. the i such that $\mathcal{D}(i) = \emptyset$. Then, problem (20) is reduced to the classical problem of computing the volume of \mathbf{U}_i , whose optimal solution is exactly $\mu_i^* = \lambda_{\mathbf{U}_i} = \mathbb{1}_{\mathbf{U}_i} \lambda^{n_i}$ (because $\mathcal{D}(i) = \emptyset \Rightarrow n_i = \dim \mathbf{X}_i$), which is the expected result.

Then we can proceed to the induction: let i be a node of T that is not a leaf: $\mathcal{D}(i) \neq \emptyset$; and suppose that for $j \in \mathcal{D}(i)$ such that $(C_i, C_j) \in \mathcal{E}$,

$$\mu_j^* = \mathbb{1}_{\mathbf{U}_j} \left(\prod_{k \in \mathcal{D}(j)} \mathbb{1}_{\mathbf{U}_k} \circ \pi_{\mathbf{X}_k} \lambda^{q_k} \right)^{\mathbf{X}_j \cap \mathbf{Z}_j} \otimes \lambda^{n_j}.$$

Then, constraint (21) can be rewritten as

$$\mu_i \leq \left(\bigotimes_{(C_i, C_j) \in \mathcal{E}} \left(\mathbb{1}_{\mathbf{U}_j} \left(\prod_{k \in \mathcal{D}(j)} \mathbb{1}_{\mathbf{U}_k} \circ \pi_{\mathbf{X}_k} \lambda^{q_k} \right)^{\mathbf{X}_j \cap \mathbf{Z}_j} \otimes \lambda^{n_j} \right)^{\mathbf{X}_i \cap \mathbf{X}_j} \right) \otimes \lambda^{n_i}$$

which in turn is simplified into

$$\mu_i \leq \left(\bigotimes_{(C_i, C_j) \in \mathcal{E}} \left(\mathbb{1}_{\mathbf{U}_j} \left(\prod_{k \in \mathcal{D}(j)} \mathbb{1}_{\mathbf{U}_k} \circ \pi_{\mathbf{X}_k} \right) \lambda^{n_j + q_j} \right)^{\mathbf{X}_i \cap \mathbf{Z}_j} \right) \otimes \lambda^{n_i}$$

since the CIP ensures that $\mathbf{X}_i \cap \mathbf{X}_j \cap \mathbf{Z}_j = \mathbf{X}_i \cap \mathbf{Z}_j$: indeed C_j is on the path between C_i and any C_k with $k \in \mathcal{D}(j)$, so that $C_i \cap C_k \subset C_j$ and thus $\mathbf{X}_i \cap \mathbf{X}_k \subset \mathbf{X}_j$, yielding $\mathbf{X}_i \cap \mathbf{Z}_j \subset \mathbf{X}_j$. At this point one can notice that $n_j + q_j = \dim(\mathbf{X}_j + \mathbf{Z}_j)$.

Then, using the DIP, we know that if $(C_i, C_j), (C_i, C_k) \in \mathcal{E}$ with $j \neq k$ then $C_j \cap C_k = \emptyset$ and with the CIP $C_j \cap C_l = \emptyset$ for any $l \in \mathcal{D}(k)$. This yields that $(\mathbf{X}_j + \mathbf{Z}_j) \cap (\mathbf{X}_k + \mathbf{Z}_k) = \{0\}$ and thus $\mathbf{Z}_i = \bigoplus_{(C_i, C_j) \in \mathcal{E}} (\mathbf{X}_j + \mathbf{Z}_j)$, allowing to rewrite constraint (21) as

$$\mu_i \leq \left(\bigotimes_{(C_i, C_j) \in \mathcal{E}} \mathbb{1}_{\mathbf{U}_j} \left(\prod_{k \in \mathcal{D}(j)} \mathbb{1}_{\mathbf{U}_k} \circ \pi_{\mathbf{X}_k} \right) \lambda^{n_j + q_j} \right)^{\mathbf{X}_i \cap \mathbf{Z}_i} \otimes \lambda^{n_i},$$

which simplifies into

$$\mu_i \leq \left(\prod_{j \in \mathcal{D}(i)} \mathbb{1}_{\mathbf{U}_j} \circ \pi_{\mathbf{X}_j} \lambda^{q_i} \right)^{\mathbf{X}_i \cap \mathbf{Z}_i} \otimes \lambda^{n_i}.$$

Eventually, we are again faced to a classical instance of the dense volume problem for \mathbf{U}_i , with $\left(\prod_{j \in \mathcal{D}(i)} \mathbb{1}_{\mathbf{U}_j} \circ \pi_{\mathbf{X}_j} \lambda^{q_i} \right)^{\mathbf{X}_i \cap \mathbf{Z}_i} \otimes \lambda^{n_i}$ instead of only the uniform Lebesgue measure, and we know that the optimal solution is obtained by multiplying this non-negative dominating measure with the indicator of \mathbf{U}_i , yielding

$$\mu_i^* = \mathbb{1}_{\mathbf{U}_i} \left(\prod_{j \in \mathcal{D}(i)} \mathbb{1}_{\mathbf{U}_j} \circ \pi_{\mathbf{X}_j} \lambda^{q_i} \right)^{\mathbf{X}_i \cap \mathbf{Z}_i} \otimes \lambda^{n_i}$$

which is the announced result.

We conclude by using the fact that $\mathcal{D}(1) = \{2, \dots, m\}$ and $\mathbb{R}^n = \mathbf{X}_1 + \mathbf{Z}_1 = \mathbf{X}_1 \oplus (\mathbf{X}_1^\perp \cap \mathbf{Z}_1)$ to compute the value:

$$\begin{aligned} \int_{\mathbf{X}_1} d\mu_1^* &= \int_{\mathbf{X}_1} \mathbb{1}_{\mathbf{U}_1} \left(\prod_{j=2}^m \mathbb{1}_{\mathbf{U}_j} \circ \pi_{\mathbf{X}_j} \lambda^{q_1} \right)^{\mathbf{X}_1 \cap \mathbf{Z}_1} d\lambda^{n_1} \\ &= \int_{\mathbf{X}_1} \mathbb{1}_{\mathbf{U}_1}(\mathbf{x}_1) \left(\int_{\mathbf{X}_1^\perp \cap \mathbf{Z}_1} \prod_{j=2}^m \mathbb{1}_{\mathbf{U}_j} \circ \pi_{\mathbf{X}_j}(\mathbf{x}_1 + \mathbf{z}_1) d\mathbf{z}_1 \right) d\mathbf{x}_1 \\ &= \int_{\mathbb{R}^n} \left(\prod_{i=1}^m \mathbb{1}_{\mathbf{U}_i} \circ \pi_{\mathbf{X}_i}(\mathbf{x}) \right) d\mathbf{x} \\ &= \int_{\mathbb{R}^n} \left(\prod_{i=1}^m \mathbb{1}_{\mathbf{K}_i}(\mathbf{x}) \right) d\mathbf{x} \\ &= \int_{\mathbb{R}^n} \mathbb{1}_{\mathbf{K}}(\mathbf{x}) d\mathbf{x} \\ &= \text{vol } \mathbf{K}. \end{aligned}$$

□

Therefore one obtains a sequence of infinite dimensional LPs on measures that can be algorithmically addressed using the usual SDP relaxations. The computations start from the leaves of the clique tree and proceed down to the root. It is worth noting that all the maximal cliques of the same generation in the tree are totally independent, which allows to treat them simultaneously, *i.e.* to partially parallelize the computations. Let $d \in \mathbb{N}$. We consider the solutions $\mathbf{m}_i^{(d)}$ to the moment relaxations, for $i = 1, \dots, m$:

$$\mathbf{P}_{d,i} = \max_{\mathbf{m}_i, \hat{\mathbf{m}}_i \in \mathbb{R}^{s(d)}} m_{i,0} \quad (23)$$

$$\text{s.t.} \quad m_{i,(\alpha^{(j)})_j, \beta} + \hat{m}_{i,(\alpha^{(j)})_j, \beta} = \left(\prod_{(C_i, C_j) \in \mathcal{E}} m_{j, \alpha^{(j)}, 0}^{(d)} \right) \ell_{i, \beta} \quad (24)$$

$$\mathbf{M}_d(\mathbf{m}_i) \succeq 0, \mathbf{M}_d(\hat{\mathbf{m}}_i) \succeq 0$$

$$\mathbf{M}_{d-d_i}(\mathbf{g}_i \mathbf{m}_i) \succeq 0$$

where $(\alpha^{(j)})_j, \beta$ are appropriate multi-indices and $\ell_{i, \beta}$ is the β moment of λ^{n_i} on the appropriate projection of \mathbf{B} . We are going to study the convergence of the sequence $\mathbf{P}_{d,1}$ to $\text{vol } \mathbf{K}$.

Theorem 19 (Convergence of the branched Moment-SOS hierarchy). *Let $C_i \in \mathcal{K}$ such that for $C_j \in \mathcal{K}$ satisfying $(C_i, C_j) \in \mathcal{E}$, we have a converging sequence of moment vectors $(\mathbf{m}_j^{(d)})_d$: $\mathbf{m}_j^{(d)} \in \mathbb{R}^{s(d)}$ and for any appropriate multi-index α , $m_{j, \alpha}^{(d)} \xrightarrow{d \rightarrow \infty} \int \mathbf{x}^\alpha d\mu_j^*$.*

In this setting, if each relaxation of the LP problem (23) associated to the clique C_i has at least one feasible solution, then the Moment-SOS hierarchy associated to clique C_i converges, in the sense that for any appropriate multi-index α , $m_{i, \alpha}^{(d)} \xrightarrow{d \rightarrow \infty} \int \mathbf{x}^\alpha d\mu_i^$.*

Thus, by induction, if at all nodes of T the moment relaxations remain feasible at all degrees of relaxation, then the branched Moment-SOS hierarchy converges, namely $\mathbf{P}_{d,1} \xrightarrow{d \rightarrow \infty} \text{vol } \mathbf{K}$.

Proof. The feasibility assumption ensures that the $\mathbf{m}_i^{(d)}$ are properly defined at all degrees d . Then, pointwise convergence of the $(\mathbf{m}_j^{(d)})_d$ yields that $m_{j,0}^{(d)}$ is bounded for the weak-* topology on $\mathbb{R}[\mathbf{x}]'$. Constraint (24) yields that $(\mathbf{m}_i^{(d)})_d$ is bounded for the weak-* topology on $\mathbb{R}[\mathbf{x}]'$, which means, according to Banach-Alaoglu's theorem, that it has an accumulation point. Finally, by uniqueness of the solution to the infinite dimensional LP problem (20), the convergence of $(\mathbf{m}_j^{(d)})_d$ to the moment sequence of μ_j^* ensures that this accumulation point is none other than the moment sequence of μ_i^* . This proves existence and uniqueness of the accumulation point of $(\mathbf{m}_i^{(d)})_d$. Then, we get for any appropriate multi-index α

$$m_{i, \alpha}^{(d)} \xrightarrow{d \rightarrow \infty} \int \mathbf{x}^\alpha d\mu_i^*.$$

To conclude for the global convergence of the sparse scheme, we just need to check the base case of this induction. Here again the base case is the leaves of the tree at which we are face to standard instances of the volume problem, whose associated Moment-SOS hierarchy is already proved to converge. Thus, our convergence assumption is satisfied, which means that as long as all the relaxations are feasible, their solutions converge weakly-* to the infinite dimensional optimal measures, and in particular

$$\mathbf{P}_{d,1} \xrightarrow{d \rightarrow \infty} \text{vol } \mathbf{K}.$$

□


Remark 20. Stokes constraints can be implemented similarly to the linear case.

6.3 Example: 6D polytope


Let $\mathbf{X} := \mathbb{R}^6$ and $\mathbf{X}_1 = \langle x_1, x_2 \rangle$, $\mathbf{X}_2 = \langle x_2, x_3, x_4 \rangle$, $\mathbf{X}_3 = \langle x_3, x_5 \rangle$, $\mathbf{X}_4 = \langle x_4, x_6 \rangle$. For $i = 1, 3, 4$ let $\mathbf{g}_i(x, y) := (x, y, 1 - x - y)$ and $\mathbf{U}_i := \mathbf{g}_i^{-1}(\mathbb{R}_+^2) = \{(u, v) \in [0, 1]^2 : u + v \leq 1\}$. Let $\mathbf{g}_2(x, y, z) := (x, y, z, 1 - x - y - z)$ and $\mathbf{U}_2 := \mathbf{g}_2^{-1}(\mathbb{R}_+^3) = \{(x, y, z) \in [0, 1]^3 \mid x + y + z \leq 1\}$. Let us approximate the volume of the 6D polytope

$$\mathbf{K} := \{\mathbf{x} \in \mathbb{R}_+^6 : x_1 + x_2 \leq 1, x_2 + x_3 + x_4 \leq 1, x_3 + x_5 \leq 1, x_4 + x_6 \leq 1\} = \bigcap_{i=1}^4 \pi_{\mathbf{X}_i}^{-1}(\mathbf{U}_i).$$

No linear clique tree is associated to this problem through Proposition 16. The only possible clique trees for applying our method are the two branched clique trees of Figure 11.



(a) 3 step clique tree



(b) 2 step clique tree

Figure 11: Two possible branched clique trees for the 6D polytope.

Let us compare the performance of the algorithms derived from the two possible clique tree configurations and with the dense problem. For that, we first write the problem associated with the 3 step clique tree configuration of the top of Figure 11:

$$\text{vol } \mathbf{K} = \int_{\mathbf{X}_1} d\mu_1^* \quad (25)$$

where

$$\begin{aligned}
\mu_1^* &= \operatorname{argmax}_{\mu_1 \in \mathcal{M}_+(\mathbf{X}_1)} \int_{\mathbf{X}_1} d\mu_1 \\
\text{s.t. } & d\mu_1(x_1, x_2) \leq dx_1 d\mu_2^{*(x_2)}(x_2) \\
& \text{spt } \mu_1 \subset \mathbf{U}_1 \\
\mu_2^* &= \operatorname{argmax}_{\mu_2 \in \mathcal{M}_+(\mathbf{X}_2)} \int_{\mathbf{X}_2} d\mu_2 \\
\text{s.t. } & d\mu_2(x_2, x_3, x_4) \leq dx_2 d\mu_3^{*(x_3)}(x_3) d\mu_4^{*(x_4)}(x_4) \\
& \text{spt } \mu_2 \subset \mathbf{U}_2 \\
\mu_i^* &= \operatorname{argmax}_{\substack{\mu_i \in \mathcal{M}_+(\mathbf{X}_i) \\ i=3,4}} \int_{\mathbf{X}_i} d\mu_i \\
\text{s.t. } & \mu_i \leq \lambda^2 \\
& \text{spt } \mu_i \subset \mathbf{U}_i, \quad i = 3, 4.
\end{aligned}$$

This problem can be complemented with the following Stokes constraints:

$$\begin{aligned}
\frac{\partial}{\partial x_1} [x_1 (1 - x_1 - x_2)] d\mu_1(x_1, x_2) &= \frac{\partial}{\partial x_1} [x_1 (1 - x_1 - x_2) d\mu_1(x_1, x_2)] \\
\frac{\partial}{\partial x_2} [x_2 (1 - x_2 - x_3 - x_4)] d\mu_2(x_2, x_3, x_4) &= \frac{\partial}{\partial x_2} [x_2 (1 - x_2 - x_3 - x_4) d\mu_2(x_2, x_3, x_4)] \\
\frac{\partial}{\partial x_i} [x_i (1 - x_i - x_{i+2})] d\mu_i(x_i, x_{i+2}) &= \frac{\partial}{\partial x_i} [x_i (1 - x_i - x_{i+2}) d\mu_i(x_i, x_{i+2})] \\
\frac{\partial}{\partial x_{i+2}} [x_{i+2} (1 - x_i - x_{i+2})] d\mu_i(x_i, x_{i+2}) &= \frac{\partial}{\partial x_{i+2}} [x_{i+2} (1 - x_i - x_{i+2}) d\mu_i(x_i, x_{i+2})] \quad i = 3, 4.
\end{aligned}$$

The 2 step clique tree of the bottom of Figure 11 yields the following formulation

$$\operatorname{vol} \mathbf{K} = \int_{\mathbf{X}_2} d\mu_2^* \tag{26}$$

where

$$\begin{aligned}
\mu_2^* &= \operatorname{argmax}_{\mu_2 \in \mathcal{M}_+(\mathbf{X}_2)} \int_{\mathbf{X}_2} d\mu_2 \\
\text{s.t. } & d\mu_2(x_2, x_3, x_4) \leq d\mu_1^{*(x_2)}(x_2) d\mu_3^{*(x_3)}(x_3) d\mu_4^{*(x_4)}(x_4) \\
& \text{spt } \mu_2 \subset \mathbf{U}_2 \\
\mu_i^* &= \operatorname{argmax}_{\substack{\mu_i \in \mathcal{M}_+(\mathbf{X}_i) \\ i=1,3,4}} \int_{\mathbf{X}_i} d\mu_i \\
\text{s.t. } & \mu_i \leq \lambda_{\mathbf{X}_i} \\
& \text{spt } \mu_i \subset \mathbf{U}_i, \quad i = 1, 3, 4
\end{aligned}$$

with Stokes constraints

$$\begin{aligned}
\frac{\partial}{\partial x_i} [x_i (1 - x_i - x_j)] d\mu_i(x_i, x_j) &= \frac{\partial}{\partial x_i} [x_i (1 - x_i - x_j) d\mu_i(x_i, x_j)] \\
\frac{\partial}{\partial x_j} [x_j (1 - x_i - x_j)] d\mu_i(x_i, x_j) &= \frac{\partial}{\partial x_j} [x_j (1 - x_i - x_j) d\mu_i(x_i, x_j)]
\end{aligned}$$

for $(i, j) = (1, 2), (3, 5), (4, 6)$. However, no Stokes constraints can be applied for the computation of μ_2 (there is no Lebesgue measure in the domination constraint, so the optimal

measure is not uniform). For this reason one can expect a slower convergence than in the linear configuration.

We implement the hierarchies associated to the 2 and 3 step sparse formulations, as well as the dense problem hierarchy, and compare their performance in Figure 12. We can compute analytically

$$\text{vol } \mathbf{K} = \frac{1}{18} \simeq 0.0556.$$

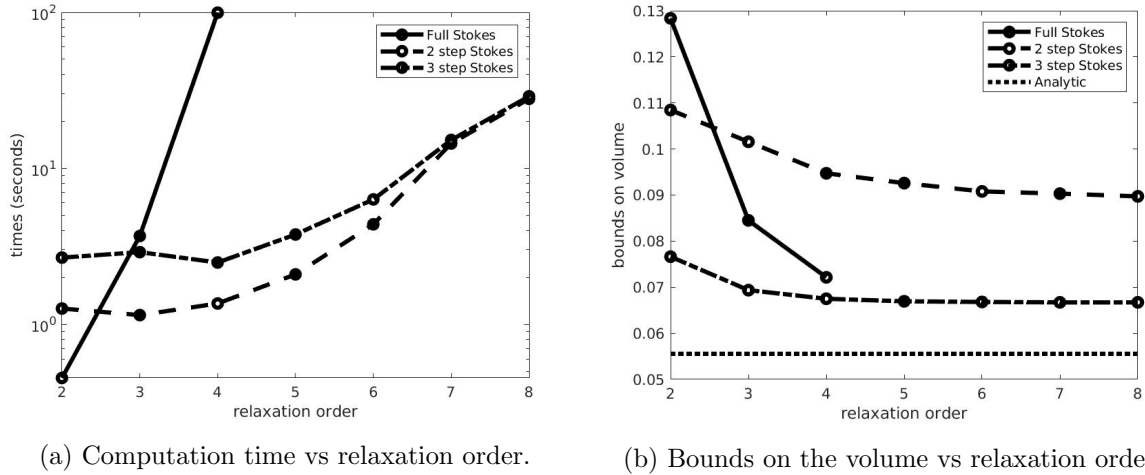


Figure 12: Performance for the 6D polytope.

Both sparse formulations outperform the dense one in terms of computational time needed to solve the corresponding SDPs (Figure 12a). On the accuracy side however (Figure 12b), we observe that the 2 step formulation seems to be less efficient than the 3 step formulation. In particular when considering the accuracy/time effort relation at order 3 the dense formulation provides a better value in almost the same time.

We believe that this can be explained by the way Stokes constraints are added to the program. Indeed, at a given clique, Stokes constraints can only be implemented in the variables that are not shared with the input measure. In the fully (2 step) branched configuration, the last step of the optimization program cannot be accelerated by Stokes constraints at all, while in the 3 step configuration, step 1 includes Stokes constraints in x_3, x_4, x_5, x_6 , step 2 includes Stokes constraints in x_2 and step 3 includes Stokes constraints on x_1 , which explains the gap between the optimal values of these two configurations.

Moreover, it seems that even the least branched (3 step) configuration still presents a gap between its optimal value and the analytic solution. This might also happen with a non-sparse instance of the Moment-SOS hierarchy (which converges theoretically) and it is likely due to the choice of the monomial basis to represent polynomials. Indeed, most of the Moment-SOS parsers generate SDP problems with the basis of monomials, while sometimes other bases (e.g. Chebyshev or Legendre polynomials) are more appropriate. However, in this precise case, it might also be linked again with the sparse Stokes constraints implementation. Indeed, in step 2 of the scheme, the unknown measure measures x_2, x_3 and x_4 but Stokes constraints are implemented only in x_2 , leaving a possible Gibbs effect in x_3, x_4 . Unlike most of our numerical examples, this one still includes an optimization step in which most of the variables are not controlled through Stokes constraints. The gap between the optimal value and the analytic value for the 3 step branched formulation in Figure 12b could be explained by a Gibbs effect in the second optimization step.

As a consequence, in the following, one should avoid the branched hierarchies that cannot be accelerated at each step at least partially by Stokes constraints. Such a hierarchy appears

when the root of the chosen clique tree shares all its vertices with its children cliques. It can be proved that such a configuration can always be avoided while implementing sparse volume computation, by choosing a leaf as the new root of the tree.

6.4 Example: 4D polytope

Let $\mathbf{X} := \mathbb{R}^4$, $\mathbf{X}_1 := \langle x_1, x_2 \rangle$, $\mathbf{X}_2 := \langle x_2, x_3 \rangle$, $\mathbf{X}_3 := \langle x_3, x_4 \rangle$, $\mathbf{g}_i(u, v) := (u, v, 1 - u - v)$, $i = 1, 2, 3$ and $\mathbf{U}_i := \mathbf{g}_i^{-1}(\mathbb{R}_+^3) = \{(u, v) \in [0, 1]^2 : u + v \leq 1\}$. Let us approximate the volume of the 4D polytope

$$\mathbf{K} := \{(x_1, x_2, x_3, x_4) \in \mathbb{R}_+^4 : x_1 + x_2 \leq 1, x_2 + x_3 \leq 1, x_3 + x_4 \leq 1\} = \bigcap_{i=1}^3 \pi_{\mathbf{X}_i}^{-1}(\mathbf{U}_i).$$

In such a case, there are two possible configurations for the associated clique tree of Proposition 16, see Figure 13. Accordingly, we can compute $\text{vol } \mathbf{K}$ in two different ways. The first way




Figure 13: Two possible clique trees for the 4D polytope.

$$\begin{aligned} \text{vol } \mathbf{K} &= \max_{\substack{\mu_1 \in \mathcal{M}_+(\mathbf{X}_1) \\ \mu_2 \in \mathcal{M}_+(\mathbf{X}_2) \\ \mu_3 \in \mathcal{M}_+(\mathbf{X}_3)}} \int_{\mathbf{X}_1} d\mu_1 & (27) \\ \text{s.t. } & d\mu_1(x_1, x_2) \leq dx_1 d\mu_2^{(x_2)}(x_2) \\ & d\mu_2(x_2, x_3) \leq dx_2 d\mu_3^{(x_3)}(x_3) \\ & d\mu_3(x_3, x_4) \leq dx_3 dx_4 \\ & \text{spt } \mu_i \subset \mathbf{U}_i, \quad i = 1, 2, 3. \end{aligned}$$

is the linear formulation given by Corollary 8, which is under the form of a non-parallelizable single linear problem. The following additional Stokes constraints can be added:

$$\begin{aligned} \frac{\partial}{\partial x_1} [x_1 (1 - x_1 - x_2)] d\mu_1(x_1, x_2) &= \frac{\partial}{\partial x_1} [x_1 (1 - x_1 - x_2) d\mu_1(x_1, x_2)] \\ \frac{\partial}{\partial x_2} [x_2 (1 - x_2 - x_3)] d\mu_2(x_2, x_3) &= \frac{\partial}{\partial x_2} [x_2 (1 - x_2 - x_3) d\mu_2(x_2, x_3)] \\ \frac{\partial}{\partial x_3} [x_3 (1 - x_3 - x_4)] d\mu_3(x_3, x_4) &= \frac{\partial}{\partial x_3} [x_3 (1 - x_3 - x_4) d\mu_3(x_3, x_4)] \end{aligned}$$

$$\frac{\partial}{\partial x_4} \left[x_4 (1 - x_3 - x_4) \right] d\mu_3(x_3, x_4) = \frac{\partial}{\partial x_4} [x_4 (1 - x_3 - x_4) d\mu_3(x_3, x_4)].$$

On the other hand, if one associates the maximal clique C_1 to the subspace \mathbf{X}_2 and the maximal clique C_2 to the subspace \mathbf{X}_1 , one also has

$$\text{vol } \mathbf{K} = \int_{\mathbf{X}_2} d\mu_2^* \quad (28)$$

where

$$\begin{aligned} \mu_2^* &= \operatorname{argmax}_{\mu_2 \in \mathcal{M}_+(\mathbf{X}_2)} \int_{\mathbf{X}_2} d\mu_2 \\ \text{s.t. } & d\mu_2(x_2, x_3) \leq \mu_1^{*(x_2)}(dx_2) d\mu_3^{*(x_3)}(x_3) \\ & \text{spt } \mu_2 \subset \mathbf{U}_2 \\ \mu_1^* &= \operatorname{argmax}_{\mu_1 \in \mathcal{M}_+(\mathbf{X}_1)} \int_{\mathbf{X}_1} d\mu_1 \\ \text{s.t. } & d\mu_1(x_1, x_2) \leq dx_1 dx_2 \\ & \text{spt } \mu_1 \subset \mathbf{U}_1 \\ \mu_3^* &= \operatorname{argmax}_{\mu_3 \in \mathcal{M}_+(\mathbf{X}_3)} \int_{\mathbf{X}_3} d\mu_3 \\ \text{s.t. } & d\mu_3(x_3, x_4) \leq dx_3 dx_4 \\ & \text{spt } \mu_3 \subset \mathbf{U}_3 \end{aligned}$$

which is the branched formulation associated to Theorem 18. Here one can see that μ_1^* and μ_3^* can be computed independently in parallel, and then re-injected in the problem to which μ_2^* is the solution. One can add the following Stokes constraints:

$$\begin{aligned} \frac{\partial}{\partial x_1} \left[x_1 (1 - x_1 - x_2) \right] d\mu_1(x_1, x_2) &= \frac{\partial}{\partial x_1} [x_1 (1 - x_1 - x_2) d\mu_1(x_1, x_2)] \\ \frac{\partial}{\partial x_2} \left[x_2 (1 - x_1 - x_2) \right] d\mu_1(x_1, x_2) &= \frac{\partial}{\partial x_2} [x_2 (1 - x_1 - x_2) d\mu_1(x_1, x_2)] \\ \frac{\partial}{\partial x_3} \left[x_3 (1 - x_3 - x_4) \right] d\mu_3(x_3, x_4) &= \frac{\partial}{\partial x_3} [x_3 (1 - x_3 - x_4) d\mu_3(x_3, x_4)] \\ \frac{\partial}{\partial x_4} \left[x_4 (1 - x_3 - x_4) \right] d\mu_3(x_3, x_4) &= \frac{\partial}{\partial x_4} [x_4 (1 - x_3 - x_4) d\mu_3(x_3, x_4)]. \end{aligned}$$

We can compute analytically

$$\text{vol } \mathbf{K} = \frac{5}{24} \simeq 0.2083.$$

In Figure 14 we compare the two sparse formulations with and without Stokes constraints. Surprisingly the linear formulation is faster than the branched one for small relaxation degrees, most probably because at this level of precision the branching costs more in terms of constructing and parsing the LMIs than it saves in computational time. When going deeper in the hierarchy we see the advantage of the branched formulation where more computations are done in parallel. As observed in the previous example however, the branched formulation seems to have problems to converge to the optimal value on an early relaxation. While the values of both formulations without Stokes constraints almost coincide, the values of the linear formulation with Stokes are strictly better than the ones of the accelerated branched formulation. This further supports our conjecture that formulations where Stokes constraints can be added at every step of the optimization program are to be preferred: the fact that both configurations behave equally

without Stokes constraints and that the branched configuration keeps a relaxation gap when implementing Stokes constraints suggests that these Stokes constraints behave better in linear configurations than in branched configurations. For this reason, in the 6D case where all possible configurations are branched, we could not completely eliminate the relaxation gap, while in this case where there is a linear configuration, the relaxation gap vanishes.




Figure 14: Performance for the 4D polytope.

6.5 Discussion on the DIP assumption

It may happen that Assumption 16 does not hold, in which case all the above results would not apply. For example one could think of the following set:

$$\mathbf{K} := \{ \mathbf{x} \in \mathbb{R}^6 : x_1^2 + x_2^2 + x_3^2 \leq 1, x_i^2 + x_{i+2}^2 + x_{i+3}^2 \leq 1, i = 2, 3 \}$$

whose correlative graph is represented on Figure 15. Here the DIP and CIP cannot be simultaneously enforced: the CIP would only be satisfied by a branched clique tree, but since all the cliques share common vertices, in such a branched tree there would automatically be sibling cliques with nonempty intersection. Also, one can notice that in this case (and, as far as we know, only in similar configurations where Assumption 16 is violated), the clique C_2 does not correspond to a polynomial appearing in the description of \mathbf{K} .

First, we would like to emphasize that the core of this article is the linear computation theorem, for which the working assumption always holds. The branched generalizations are only consequences of this linear computation theorem.

Second, the fact that Assumption 16 does not hold is not a dead-end for using our scheme. In fact, even the simpler CIP might not hold, in which case one would need to perform a *chordal extension*, which consists of adding virtual links between variables to construct a chordal graph. Basically, a chordal extension would make the graph chordal at the price of slightly weakening the correlative sparsity pattern. In general, the same can be performed to enforce Assumption 16: one could add virtual links between variables to enforce the assumption to hold. For example, if one artificially links variables x_3 and x_4 in the correlation graph of \mathbf{K} , one obtains a new correlation graph, with an associated clique tree satisfying all our working assumptions (see Figure 16). This manipulation results in increasing the correlative sparsity from 3 to 4, which is a weakening of the correlative sparsity pattern. However, our framework still allows to reduce the dimensionality of the problem from 6 to 4.

Remark 21. There might exist examples for which the DIP would only be obtained by completely destroying the sparsity pattern one wants to exploit. However, this could also happen




Figure 15: A correlation graph that violates Assumption 16.




Figure 16: A way to fix our counterexample.

with the more common CIP. In either case, an option might be to consider the dual problem of (2) as a way to find minimizing sequences $(v_k)_k$ that approximate $\mathbb{1}_{\mathbf{K}}$, and to apply it to each one of the $\mathbf{U}_i := \{\mathbf{x}_i \in \mathbf{X}_i : \mathbf{g}_i(\mathbf{x}_i) \geq 0\}$ to obtain a sequence $(v_1^{(k)}, \dots, v_m^{(k)})_k$ such that $(v_i^{(k)})_k$ approximates $\mathbb{1}_{\mathbf{U}_i}$. Then, one would still have to prove that the convergence of the hierarchy is stable by product (which is nontrivial) to conclude that $\int \prod_{i=1}^m v_i^{(k)} \circ \pi_{\mathbf{X}_i}(\mathbf{x}) d\mathbf{x}$ converges to $\text{vol } \mathbf{K}$. The major drawback of this solution is that in order for $(v_k)_k$ to approximate $\mathbb{1}_{\mathbf{K}}$, we cannot implement Stokes constraints, as they would modify problem (2) (see Remark 13) in a way that makes us lose the convergence to an indicator (this is precisely the point of Stokes constraints: they allow to obtain the volume without trying to approximate discontinuous functions with polynomials). However, we know that the volume approximation hierarchy has a bad convergence rate without Stokes constraints. In general, we should not expect any method that approximates indicators with polynomials to yield satisfactory results in terms of volume computation. Such a method should be considered only as a last resort if any trial to apply

the above scheme fails. Finally, this option represents a framework that would be completely independent from the above and thus it remains out of the scope of this paper.

Conclusion

Our results

In this paper we addressed the problem of approximating the volume of sparse semi-algebraic sets with the Moment-SOS hierarchy of SDP relaxations. As illustrated by our examples, our sparse formulation allows to dramatically decrease the computational time for each relaxation, and to tackle high dimensional volume computation problems that are not tractable with the usual SDP methods. By splitting the problems into low dimensional subproblems, one drastically reduces the dimension of each relaxation, without loss of precision. This reduction of complexity is due to the correspondance between the structure of our algorithm and the correlative sparsity pattern in the description of the semi-algebraic set.

We also showed that additional Stokes constraints have a huge effect on convergence and precision for volume computation, and that they can successfully be adapted to our sparse formulations. This yields a much better rate of convergence for the corresponding hierarchy. However, implementing these Stokes constraints leads to subtle constraints that have to be enforced if one wants to efficiently compute the volume:

- First, one should always prefer the linear formulation of Theorem 8 whenever possible, since this ensures that Stokes constraints can always be efficiently implemented.
- Then, in the more general case of Theorem 18, one should always avoid formulations in which the root at the computation tree has no Stokes constraint; fortunately, such configurations can always be avoided by choosing a leaf as the root of the clique tree.

Furthermore, in the branched case, one should be aware of the fact that each step of the algorithm introduces an approximation error, and the errors accumulate until the root is reached. Consequently, a formulation in which the clique tree has too many generations will lead to a larger global error than a formulation with less generations. For this reason, one should minimize the number of generations in the clique tree, which is equivalent to parallelizing as much as possible. In addition to that, when the problem has many dimensions and branches, parallelization can obviously drastically increase the speed of the computations.

Applications and future work

To the best of our knowledge, this sparse method for solving volume problems is new and full of promises for future applications. For instance, the problem of computing the mass of any compactly supported measure absolutely continuous (with respect to the Lebesgue measure) can be addressed using this sparsity method. Also, measures that are not compactly supported but have some decay properties (e.g. Gaussian measures) can also be handled by our method, which may prove useful in computations for probability and statistics. Also, specific constraints could probably be used in addition to Stokes constraints when the semi-algebraic set presents a specific structure (e.g. a polytope, a convex body).

Furthermore, the framework of exploiting correlative sparsity can be applied to any method that relies on computations on measures, whether these measures are represented by their moments (as it is done in this paper), or by samples (as in the stochastic volume computation methods). In particular, we believe that this formalism could easily be extended to Monte-Carlo-based volume computations.

Finally, we also believe that this method can be adapted to the computation of regions of attraction, through the formalism developed in [13], for high dimensional differential-algebraic

systems that present a network structure, such as power grids, distribution networks in general and possibly other problems. The main difficulty resides in taking sparsity into account when formulating the Liouville equation, and keeping uniqueness of the solution as in the non-controlled non-sparse framework.

Acknowledgements

We are grateful to an anonymous Reviewer for a deep insight that helped us significantly improve the initial version of this paper.

A Disjoint Intersection Property

Definition 22. Let $G = (V, E)$ be a graph with vertices set V and edges set $E \subset V^2$. The following definitions can be found in e.g. [19]:

- The *degree* $\deg v$ of $v \in V$ is the cardinal of the set $\{w \in V : (v, w) \in E\}$ i.e. the number of vertices connected to v .
- A *clique* of G is a subset of vertices $C \subset V$ such that $u, v \in C$ implies $(u, v) \in E$.
- A graph G is *chordal* if every cycle of length greater than 3 has a chord, i.e. an edge connecting two nonconsecutive vertices on the cycle.
- A *tree* $T = (\mathcal{K}, \mathcal{E})$ is a graph without cycle.
- The *treewidth* of a chordal graph is the size of its biggest clique minus 1. Thus the treewidth of a tree is 1 and the treewidth of a complete graph ($E = V^2$) of size n is $n - 1$.
- A *rooted tree* is a tree in which one vertex has been designated the *root*.
- In a rooted tree, the *parent* of a vertex v is the vertex w connected to it on the path to the root; v is then called a *child* of w ; two vertices that have the same parent are called *siblings*; a *descendant* of a vertex v is any vertex which is either the child of v or is (recursively) the descendant of any of the children of v ; v is then called an *ancestor* of itself and any of its descendants.
- The vertices of a rooted tree can be partitioned between the root, the *leaves* (the vertices that have parents but no children) and the *branches* (that have children and parents).
- Let \mathcal{K} be the set of maximal cliques of G . A *clique tree* $T = (\mathcal{K}, \mathcal{E})$ of G is a tree whose vertices are the maximal cliques of G .
- A clique tree satisfies the clique intersection property (CIP) if for every pair of distinct cliques $C, C' \in \mathcal{K}$, the set $C \cap C'$ is contained in every clique on the path connecting C and C' in the tree. We denote by \mathcal{T}^{ct} the set of clique trees of G that satisfy the CIP.
- Let $A \subset V$. Then, the *subgraph of G generated by A* is given by $\langle A \rangle_G := (A, E \cap A^2)$.

Theorem 23. A connected graph G is chordal if and only if $\mathcal{T}^{ct} \neq \emptyset$ if and only if \mathcal{K} admits an ordering that satisfies the RIP.

Definition 24. Let $G = (V, E)$ be chordal and connected. Let $T = (\mathcal{K}, \mathcal{E}) \in \mathcal{T}^{ct}$ be a clique tree rooted in $C_1 \in \mathcal{K}$. T satisfies the *Disjoint Intersection Property* (DIP) if $\forall C, C', C'' \in \mathcal{K}$, if $(C, C') \in \mathcal{E}$ and $(C, C'') \in \mathcal{E}$ then $C' = C''$ or $C' \cap C'' = \emptyset$. In words, each clique has an empty intersection with all its siblings.

We are now going to give a systematic way to enforce Assumption 16 and generate the associated clique trees. Let $(\mathbf{g}_1, \dots, \mathbf{g}_m)$ be a correlative sparse family of polynomial vectors with a connected chordal correlation graph $G = (V, E)$. Let \mathcal{K} be the set of maximal cliques of G . We construct the *clique graph* $G_{\mathcal{K}} = (\mathcal{K}, \mathcal{F})$ such that $(C, C') \in \mathcal{F}$ iff $C \cap C' \neq \emptyset$. One can in turn define cliques (called *meta-cliques*) for this new graph, and its correlative sparsity $CS_{\mathcal{K}}$ is the size of its biggest maximal meta-clique minus 1. One can note that any clique tree is a subtree of $G_{\mathcal{K}}$ including all its vertices.

Remark 25. If $G_{\mathcal{K}}$ itself is a tree (as in section 6.3), then it trivially satisfies the DIP and CIP, and Assumption 16 automatically holds.

Lemma 26. *Let $T = (\mathcal{K}, \mathcal{E})$ be a clique tree satisfying Assumption 16.*

- 1) *Let $C, C' \in \mathcal{K}$ such that $C \neq C'$ and $C \cap C' \neq \emptyset$. Then, up to permuting them, C is a descendant of C' .*
- 2) *Let K be a meta-clique. Then, $\langle K \rangle_T$ is an oriented path of T : the elements of K are ancestor to one another and each $C \in K$ has its parent in K except one of them.*

Proof.

- 1) Let C'' be the last common ancestor of C and C' , meaning that C'' is an ancestor of both C and C' but any child of C'' is the ancestor of at most one of them. Such ancestor exists since the root C_1 is a common ancestor to C and C' . Then, C'' is on the path between C and C' . Since $C \neq C'$, up to permuting them, we can suppose that $C \neq C''$.

By contradiction, we suppose that $C' \neq C''$. Then, let \hat{C} be the child of C'' that is also the ancestor of C , and \tilde{C} the child of C'' that is also the ancestor of C' . Both exist since C'' is an ancestor of C and C' and $C \neq C' \neq C''$. C'' being the latest common ancestor of C and C' , we deduce that $\hat{C} \neq \tilde{C}$ so that \hat{C} and \tilde{C} are siblings. Then, the DIP ensures that $\hat{C} \cap \tilde{C} = \emptyset$. However, \hat{C} and \tilde{C} are on the path between C and C' , so according to the CIP they both contain $C \cap C'$ which is nonempty. This is a contradiction.

- 2) According to point 1), all elements of K are descendants of one another, so that they are all on the same path in T . We only have to show that any C between two elements C', C'' of K on this path is also an element of K . Indeed, let $C''' \in K$. Then, up to a permutation on $\{C', C'', C'''\}$ we can suppose that the unoriented path includes in this order: (C', C, C'', C''') . Then, C is on the path between C' and C''' , so the CIP implies that $C \supset C' \cap C'''$ is nonempty (since C' and C''' belong to the same clique K), and then C has a nonempty intersection with C''' . This shows that C has a nonempty intersection with any element of K , which by maximality of K is the definition of $C \in K$.

□

Corollary 27. *If $G_{\mathcal{K}}$ is a complete graph (all pairs of maximal cliques have nonempty intersection as in section 6.4), then the only candidates for our clique tree are linear clique trees. In such case, Assumption 16 is equivalent to the existence of a reordering of $(\mathbf{g}_1, \dots, \mathbf{g}_m)$ such that Assumption 7 holds.*

We now give an algorithm to generate a clique tree $T = (\mathcal{K}, \mathcal{E})$ that satisfies the DIP and the CIP. In case Assumption 16 does not hold, this algorithm automatically adds edges to E until it finds an appropriate clique tree (see Algorithm 1).

Remark 28. Algorithm 1 deserves some explanations:

- Minimizing $\deg C_1$ is a way to ensure Stokes constraints will be fully implementable, in contrast to the 2 step implementation in section 6.3.

Algorithm 1: How to build an appropriate clique tree

Data: $G = (V, E)$ and its clique graph $G_{\mathcal{K}} = (\mathcal{K}, \mathcal{F})$.

Result: $T = (\mathcal{K}, \mathcal{E})$ satisfying CIP & DIP if Assumption 16 holds, $G = (V, E)$ with additional edges else.

```
1 Initialization: Choose  $C_1 \in \mathcal{K}$  with minimal degree in  $G_{\mathcal{K}}$ ;  
2   Initialize  $i = j = k = 1$ ,  $\mathcal{P}_1 := \{C_1\}$ ,  $\mathcal{E}_1 := \emptyset$ ;  
3 while  $k < |\mathcal{K}|$  do  
4   while  $\mathbb{M}_{ik} := \{K \text{ maximal meta-clique} : K \cap \mathcal{P}_k^c \neq \emptyset, C_i \in K\} \neq \emptyset$  do  
5     Choose  $K_j \in \operatorname{argmax}_{K \in \mathbb{M}_{ik}} |K|$ ;  
6     while  $K_j \cap \mathcal{P}_k^c \neq \emptyset$  do  
7        $l := \max\{r \leq k : C_r \in K_j\}$ ;  
8       if  $\exists (C, C') \in \operatorname{argmax}\{|\hat{C} \cap C_l| : (\tilde{C}, \hat{C}) \in \mathcal{P}_k \times K_j \cap \mathcal{P}_k^c, \tilde{C} \cap \hat{C} \not\subseteq C_l\}$  then  
9         if  $|C_l \cup C| > |C_l \cup C'|$  then  
10          foreach  $v \in C', w \in C_l$  do  $E \leftarrow E \cup \{(v, w), (w, v)\}$ ;  
11          return  $G = (V, E)$  ;  
12        else  
13          foreach  $v \in C, w \in C_l$  do  $E \leftarrow E \cup \{(v, w), (w, v)\}$ ;  
14          return  $G = (V, E)$  ;  
15        end  
16       else  
17         Choose  $C_{k+1} \in \operatorname{argmax}_{C \in K_j \cap \mathcal{P}_k^c} |C \cap C_l|$   
18       end  
19       if  $\exists C \in \mathcal{P}_k$  s.t.  $(C_l, C) \in \mathcal{E}_k$  &  $C \cap C_{k+1} \neq \emptyset$  then  
20         if  $|C \cup C_{k+1}| < |C_l \cup C_{k+1}| \wedge |C_l \cup C|$  then  
21           foreach  $v \in C, w \in C_{k+1}$  do  $E \leftarrow E \cup \{(v, w), (w, v)\}$ ;  
22           return  $G = (V, E)$  ;  
23         else if  $|C_l \cup C| > |C_l \cup C_{k+1}|$  then  
24           foreach  $v \in C_{k+1}, w \in C_l$  do  $E \leftarrow E \cup \{(v, w), (w, v)\}$ ;  
25           return  $G = (V, E)$  ;  
26         else  
27           foreach  $v \in C, w \in C_l$  do  $E \leftarrow E \cup \{(v, w), (w, v)\}$ ;  
28           return  $G = (V, E)$  ;  
29         end  
30       else  
31          $\mathcal{P}_{k+1} := \mathcal{P}_k \cup \{C_{k+1}\}$ ;  
32          $\mathcal{E}_{k+1} := \mathcal{E}_k \cup \{(C_l, C_{k+1})\}$ ;  
33          $k \leftarrow k + 1$  ;  
34       end  
35     end  
36      $j \leftarrow j + 1$  ;  
37   end  
38    $i \leftarrow i + 1$  ;  
39 end  
40 return  $T := (\mathcal{P}_{|\mathcal{K}|}, \mathcal{E}_{|\mathcal{K}|})$ ;
```

- Index i denotes a clique that has already been added to the tree; the algorithm adds to the tree every clique that shares vertices with C_i , and then increments i .
- Index j denotes the meta-clique in which we are working; according to Lemma 26, in such meta-clique the cliques should be added in line.
- Index k denotes the number of elements that have already been added to the tree; when k is equal to the number of cliques, our clique tree is complete.
- Index l denotes the latest clique of the meta-clique K_j that has been added to the tree; according to Lemma 26, C_l should then be the parent of C_{k+1} .
- At line 5 we maximize $|K_j|$ to favor linear configurations as they are the most compatible with Stokes constraints.
- The if loop at line 8 checks whether it is possible to add the remaining cliques of K_j to our tree without destroying the CIP.
- The if loop at line 19 checks whether the clique we want to add destroys the DIP or not.
- At line 17 we maximize $|C_{k+1} \cap C_l|$ so that it is less likely to pose problems with CIP and DIP in the future iterations.
- The if loops at lines 9 and 20 are meant to minimize the correlative sparsity of the new graph $G = (V, E)$, since it is the limiting factor for the tractability of our algorithm.

Theorem 29. *Any clique tree returned by Algorithm 1 satisfies the DIP and the CIP.*

Proof. We are going to show by induction that at any step k , the graph $T_k := (\mathcal{P}_k, \mathcal{E}_k)$ is a tree that satisfies the CIP and the DIP. First, it is trivial that $T_1 = (\{C_1\}, \emptyset)$ is a tree and satisfies the CIP and the DIP. Next, we suppose that we have constructed a tree T_k satisfying CIP and DIP through iterations of our algorithm, and that the next iteration leads us to define a $T_{k+1} := (\mathcal{P}_k \cup \{C_{k+1}\}, \mathcal{E}_k \cup \{(C_l, C_{k+1})\})$. Since the only edge we added connected C_l to a new vertex that was not in T_k , it did not introduce any cycle, thus T_{k+1} is still a tree. We are now going to check the CIP and DIP.

- Let $C, C'' \in \mathcal{P}_{k+1}; C \cap C'' \neq \emptyset$. Let $C' \in \mathcal{P}_k$ be on the path between C and C'' in T_{k+1} ($C' \neq C_{k+1}$ because C_{k+1} is on no path in T_{k+1}).
 - If $C, C'' \in \mathcal{P}_k$ then by our induction hypothesis $C \cap C'' \subset C'$.
 - Else, without loss of generality we have $C'' = C_{k+1}$ and $C \in \mathcal{P}_k$.
 - * Since we successfully passed through the if loop of line 8, we have $C \cap C_{k+1} \subset C_l$.
 - * By our induction assumption (T_k satisfies the CIP), we have $C \cap C_l \subset C'$ (because either $C' = C_l$ or C' is on the path between C and C_l , the parent of C_{k+1}).
- This yields $C' \supset C \cap C_l \supset C \cap (C \cap C_{k+1}) = C \cap C_{k+1} = C \cap C''$.

Then, T_{k+1} satisfies the CIP.

- Let $C \in \mathcal{P}_k, C', C'' \in \mathcal{P}_{k+1}$ such that $(C, C'), (C, C'') \in \mathcal{E}_{k+1}$.
 - If $C', C'' \in \mathcal{P}_k$ then by our induction hypothesis $C' \cap C'' = \emptyset$.
 - Else, without loss of generality we have $C'' = C_{k+1}, C = C_l$, and since we successfully passed through the if loop of line 19, we have $C' \cap C'' = C' \cap C_{k+1} = \emptyset$.

Then, T_{k+1} satisfies the DIP.

□

Finally, we conjecture that if Assumption 16 holds, then Algorithm 1 will directly return a clique tree satisfying the CIP and the DIP without adding any edge to G .

B On Monte Carlo simulations

We describe the very basic Monte Carlo approach used in section 5.4. Let ξ_1, \dots, ξ_N be i.i.d. samples from some law μ . In our case μ is the uniform distribution on $[0, 1]^n$. Further let f be a function from the probability space into $\{0, 1\}$. Again, in our case, f would return 1 if the sample ξ_i is in the set \mathbf{K} , and 0 else. By the strong law of large numbers

$$X_N := \frac{1}{N} \sum_{i=1}^N f(\xi_i) \rightarrow \int f d\mu \text{ for } N \rightarrow \infty.$$

This makes X_N a reasonable guess if N is large. However, X_N is still a guess, and it might happen that X_N is actually far away from the approximated quantity.

As a consequence of the Central Limit Theorem, the difference $X_N - \int f d\mu$ behaves (almost) like a normal distributed random variable with zero mean and variance σ^2/N where $\sigma^2 = \int (f - \int f d\mu)^2 d\mu$. Note that the variance σ^2 can also be estimated based on the i.i.d. samples ξ_1, \dots, ξ_N :

$$S_N^2 := \frac{1}{N-1} \sum_{i=1}^N (\xi_i - X_N)^2.$$

This allows to define a confidence interval for the approximated volume. Indeed, say we are interested in a 99%-confidence interval. Then for G a standard normal distributed random variable, we have $P(|G| < 2.58) \approx 0.99$ and consequently,

$$P\left(X_N - \frac{2.58 S_N}{\sqrt{N}} \leq \int f d\mu \leq X_N + \frac{2.58 S_N}{\sqrt{N}}\right) \approx 0.99.$$

References

- [1] B. Bollobás. Volume estimates and rapid mixing. Pages 151–180 in *Flavors of geometry*, MSRI Publ. 31, Cambridge University Press, 1997.
- [2] M. E. Dyer, A. M. Frieze. On the complexity of computing the volume of a polyhedron. *SIAM J. Comput.* 17:967–974, 1988.
- [3] G. Elekes. A geometric inequality and the complexity of measuring the volume. *Discrete Comput. Geom.* 1:289–292, 1986.
- [4] D. Henrion, J. B. Lasserre, C. Savorgnan. Approximate volume and integration for basic semialgebraic sets. *SIAM Review* 51(4):722–743, 2009.
- [5] B. Büeler, A. Enge, K. Fukuda. Exact volume computation for polytopes: a practical study. Pages 131–154 in *Polytopes: combinatorics and computation*, G. Kalai and G. M. Ziegler (Eds.), Birkhäuser, 2000.
- [6] M. E. Dyer, A. M. Frieze, R. Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies, *J. ACM* 38:1–17, 1991.
- [7] C. Belisle. Slow hit-and-run sampling, *Statist. Probab. Lett.* 47:33–43, 2000.
- [8] C. Belisle, E. Romeijn, R. L. Smith. Hit-and-run algorithms for generating multivariate distributions. *Math. Oper. Res.* 18:255–266, 1993.
- [9] R. L. Smith. Efficient Monte Carlo procedures for generating points uniformly distributed over bounded regions. *Oper. Res.* 32:1296–1308, 1984.

- [10] B. Cousins, S. Vempala. A practical volume algorithm. *Math. Program. Comput.* 8:133–160, 2016.
- [11] J. B. Lasserre. *Moments, positive polynomials and their applications*. Imperial College Press, 2010.
- [12] J. B. Lasserre. Computing Gaussian and exponential measures of semi-algebraic sets. *Adv. Appl. Math.* 91:137–163, 2017.
- [13] D. Henrion, M. Korda. Convex computation of the region of attraction of polynomial control systems. *IEEE Trans. Autom. Control* 59(2):297-312, 2014.
- [14] C. Jozs, D.K. Molzahn, M. Tacchi, S. Sojoudi. Transient stability analysis of power systems via occupation measures. *Innovative Smart Grid Technologies*, 2019.
- [15] H. Waki, S. Kim, M. Kojima, M. Muramatsu. Sums of Squares and Semidefinite Program Relaxations for Polynomial Optimization Problems with Structured Sparsity. *SIAM J. Optim.* 17(1):218–242, 2006.
- [16] D. Cifuentes, P. A. Parrilo. Exploiting Chordal Structure in Polynomial Ideals: A Gröbner Bases Approach. *SIAM J. Discrete Math.* 30(3):1534–1570, 2016.
- [17] J. B. Lasserre. Convergent SDP relaxations in polynomial optimization with sparsity. *SIAM J. Optim.* 17:822–843, 2006.
- [18] D. Henrion, J. B. Lasserre, J. Loeftberg. GloptiPoly 3: moments, optimization and semidefinite programming. *Optimization Methods and Software* 24(4-5):761-779, 2009.
- [19] J. R. S. Blair, B. Peyton. An introduction to chordal graphs and clique trees. Pages 1–29 in *Graph Theory and Sparse Matrix Computation*, Springer, 1993.
- [20] J. B. Lasserre. Representation of chance-constraints with strong asymptotic guarantees. *IEEE Control Systems Letters* 1(1):50–55, 2017.
- [21] J. H. Hubbard, B. Burke Hubbard. *Vector calculus, linear algebra, and differential forms - A unified approach*. 2nd Ed., Prentice Hall, 2002.
- [22] R. Stanley, I. G. Macdonald, R. B. Nelsen. Solution of elementary problem E2701. *The American Mathematical Monthly* 86(5):396, 1979.
- [23] C. Jozs, D. Henrion. Strong duality in Lasserre’s hierarchy for polynomial optimization. *Optim. Lett.* 10:3-10, 2016.