

MIT Open Access Articles

Chordal Networks of Polynomial Ideals

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Cifuentes, Diego, and Pablo A. Parrilo. "Chordal Networks of Polynomial Ideals." SIAM Journal on Applied Algebra and Geometry 1, no. 1 (January 2017): 73–110. © 2017 Society for Industrial and Applied Mathematics

As Published: 10.1137/16m106995x

Publisher: Society for Industrial & Applied Mathematics (SIAM)

Persistent URL: <https://hdl.handle.net/1721.1/121461>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Chordal Networks of Polynomial Ideals*

Diego Cifuentes[†] and Pablo A. Parrilo[†]

Abstract. We introduce a novel representation of structured polynomial ideals, which we refer to as *chordal networks*. The sparsity structure of a polynomial system is often described by a graph that captures the interactions among the variables. Chordal networks provide a computationally convenient decomposition into simpler (triangular) polynomial sets, while preserving the underlying graphical structure. We show that many interesting families of polynomial ideals admit compact chordal network representations (of size linear in the number of variables), even though the number of components is exponentially large. Chordal networks can be computed for arbitrary polynomial systems using a refinement of the chordal elimination algorithm from [D. Cifuentes and P. A. Parrilo, *SIAM J. Discrete Math.*, 30 (2016), pp. 1534–1570]. Furthermore, they can be effectively used to obtain several properties of the variety, such as its dimension, cardinality, and equidimensional components, as well as an efficient probabilistic test for radical ideal membership. We apply our methods to examples from algebraic statistics and vector addition systems; for these instances, algorithms based on chordal networks outperform existing techniques by orders of magnitude.

Key words. chordal graphs, structured polynomials, chordal networks, triangular sets

AMS subject classifications. 13P15, 14Q99, 68W30

DOI. 10.1137/16M106995X

1. Introduction. Systems of polynomial equations can be used to model a large variety of applications, and in most cases the resulting systems have a particular sparsity structure. We describe this sparsity structure using a graph. A natural question that arises is whether this graphical structure can be effectively used to solve the system. In [9] we introduced the *chordal elimination* algorithm, an elimination method that always preserves the graphical structure of the system. In this paper we refine this algorithm to compute a new representation of the polynomial system that we call a *chordal network*.

Chordal networks attempt to fix an intrinsic issue of Gröbner bases: they destroy the graphical structure of the system [9, Ex. 1.2]. As a consequence, polynomial systems with simple structure may have overly complicated Gröbner bases (see [Example 1.1](#)). In contrast, chordal networks will always preserve the underlying chordal graph. We remark that chordal graphs have been successfully used in several other areas, such as numerical linear algebra [28], discrete and continuous optimization [5, 32], graphical models [23], and constraint satisfaction [11].

Chordal networks describe a decomposition of the polynomial ideal into simpler (triangular) polynomial sets. This decomposition gives quite a rich description of the underlying

*Received by the editors April 11, 2016; accepted for publication (in revised form) November 2, 2016; published electronically February 1, 2017.

<http://www.siam.org/journals/siaga/1/M106995.html>

Funding: This research was supported by AFOSR under grant FA9550-11-1-0305.

[†]Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology, Cambridge, MA 02139 (diegcif@mit.edu, parrilo@mit.edu).

variety. In particular, chordal networks can be efficiently used to compute dimension, cardinality, and equidimensional components and also to test radical ideal membership. Remarkably, several families of polynomial ideals (with exponentially large Gröbner bases) admit a compact chordal network representation of size proportional to the number of variables. We will shortly present some motivational examples after setting up the main terminology.

Throughout this document we work in the polynomial ring $\mathbb{K}[X] = \mathbb{K}[x_0, x_1, \dots, x_{n-1}]$ over some field \mathbb{K} . We fix once and for all the ordering of the variables $x_0 > x_1 > \dots > x_{n-1}$.¹ We consider a system of polynomials $F = \{f_1, f_2, \dots, f_m\}$. There is a natural graph $\mathcal{G}(F)$, with vertex set $X = \{x_0, \dots, x_{n-1}\}$, that abstracts the *sparsity structure* of F . The graph is given by cliques: for each f_i we add a clique in all its variables. Equivalently, there is an edge between x_i and x_j if and only if there is some polynomial in F that contains both variables. We will consider throughout the paper a chordal completion G of the graph $\mathcal{G}(F)$, and we will assume that $x_0 > \dots > x_{n-1}$ is a perfect elimination ordering of G (see [Definition 2.1](#)).

Some motivating examples. The notions of chordality and treewidth are ubiquitous in applied mathematics and computer science. In particular, several hard combinatorial problems can be solved efficiently on graphs of small treewidth by using some type of recursion (or dynamic program) [5]. We will see that this recursive nature is also present in several polynomial systems of small treewidth. We now illustrate this with three simple examples.

Example 1.1 (coloring a cycle graph). Graph coloring is a classical NP-complete problem that can be solved efficiently on graphs of small treewidth. We consider the cycle graph C_n with vertices $0, 1, \dots, n-1$, whose treewidth is two. Coloring C_n is particularly simple by proceeding in a recursive manner: color vertex $n-1$ arbitrarily, and then subsequently color vertex i , avoiding the color of $i+1$ and possibly $n-1$.

The q -coloring problem for a graph $\mathcal{G} = (V, E)$ can be encoded in a system of polynomial equations (see, e.g., [10]):

$$\begin{aligned} (1a) \quad & x_i^q - 1 = 0, & i \in V, \\ (1b) \quad & x_i^{q-1} + x_i^{q-2}x_j + \dots + x_ix_j^{q-2} + x_j^{q-1} = 0, & ij \in E. \end{aligned}$$

Let $F_{n,q}$ denote such a system of polynomials for the cycle graph C_n . Given that coloring the cycle graph is so easy, it should be possible to solve these equations efficiently. However, if we compute a Gröbner basis, the result is not so simple. In particular, for the case of $F_{9,3}$ one of these polynomials has 81 terms (with both lex and grevlex order). This is a consequence of the fact that Gröbner bases destroy the graphical structure of the equations.

Nonetheless, one may hope to give a simple representation of the above polynomials that takes into account their recursive nature. Indeed, a triangular decomposition of these equations is presented in [Figure 1](#) for the case of $F_{9,3}$, and the pattern is very similar for arbitrary values of n, q . The decomposition represented is

$$\mathbf{V}(F_{9,3}) = \bigcup_T \mathbf{V}(T),$$

¹Observe that smaller indices correspond to larger variables.

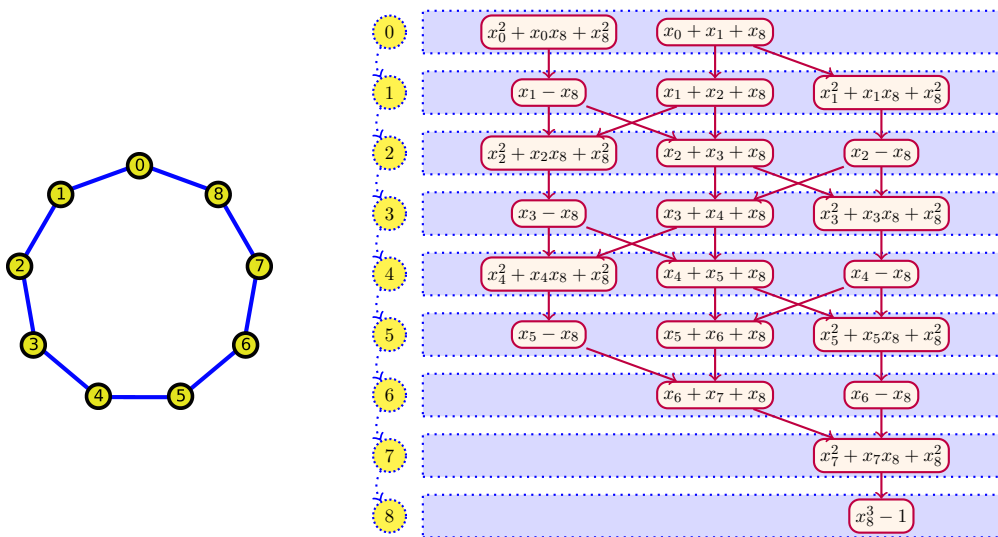


Figure 1. Chordal network for the 3-chromatic ideal of a cycle.

where the union is over all maximal directed paths in the diagram of Figure 1. One path is

$$T = \{x_0 + x_1 + x_8, x_1^2 + x_1x_8 + x_8^2, x_2 - x_8, x_3^2 + x_3x_8 + x_8^2, x_4 - x_8, x_5^2 + x_5x_8 + x_8^2, x_6 - x_8, x_7^2 + x_7x_8 + x_8^2, x_8^3 - 1\}.$$

Recall that a set of polynomials is triangular if the largest variables of these polynomials are all distinct, and observe that all maximal paths T are triangular. Note that the total number of triangular sets is 21, and in general we get the $(n - 1)$ st Fibonacci number. Even though the size of the triangular decomposition grows rapidly, it admits a very compact representation (linear in n), and the reason is precisely the recursive nature of the equations. Indeed, the diagram of Figure 1 is constructed in a way very similar to how we construct colorings: choose x_8 arbitrarily; then for each x_i choose it based on the values of x_{i+1} and x_8 .

Example 1.2 (vertex covers of a tree). We now consider the problem of finding minimum vertex coverings of a graph. Recall that a subset S of vertices is a cover if any edge is incident to at least one element in S . Since the complement of a vertex cover is an independent set, computing a minimum vertex cover is NP-complete. Nevertheless, when the graph is a tree, the minimal vertex covers have a very special structure. Indeed, we can construct such a cover recursively, starting from the root, as follows. For the root node, we can decide whether to include it in the cover or not. If we include it, we can delete the root and then recurse on each of its children. Otherwise, we need to include in the cover all of its children, so we can delete them all, and then recurse.

The minimal vertex covers of a graph $\mathcal{G} = (V, E)$ are in correspondence with the irreducible components of its *edge ideal* $I(\mathcal{G}) := \langle x_ix_j : ij \in E \rangle$ (see, e.g., [33, Prop. 7.2.3]). Therefore, the irreducible components of the edge ideal of a tree always have a very simple structure (although there might be exponentially many). For instance, the diagram in Figure 2 represents the components for the case of a simple 10-vertex tree. Here the components are given by the




Figure 2. Chordal network for the edge ideal of a tree.

possible choices of one node from each of the (purple) boxes so that these nodes are connected (e.g., $T = \{0, 0, 0, x_3, x_4, x_5, x_6, 0, x_8, 0\}$). Note that there are $2^4 + 1 = 17$ components.

Example 1.3 (adjacent minors). Let X be a $2 \times n$ matrix of indeterminates, and consider the polynomial set F_n given by its adjacent minors, i.e.,

$$X := \begin{pmatrix} x_0 & x_2 & \cdots & x_{2n-2} \\ x_1 & x_3 & \cdots & x_{2n-1} \end{pmatrix}, \quad F_n := \{x_{2i}x_{2i+3} - x_{2i+1}x_{2i+2} : 0 \leq i < n - 1\}.$$

The corresponding ideal has been studied in, e.g., [13,16]. Figure 3 shows the graph associated to this system. We are interested in describing the irreducible components of $\mathbf{V}(F_n)$.

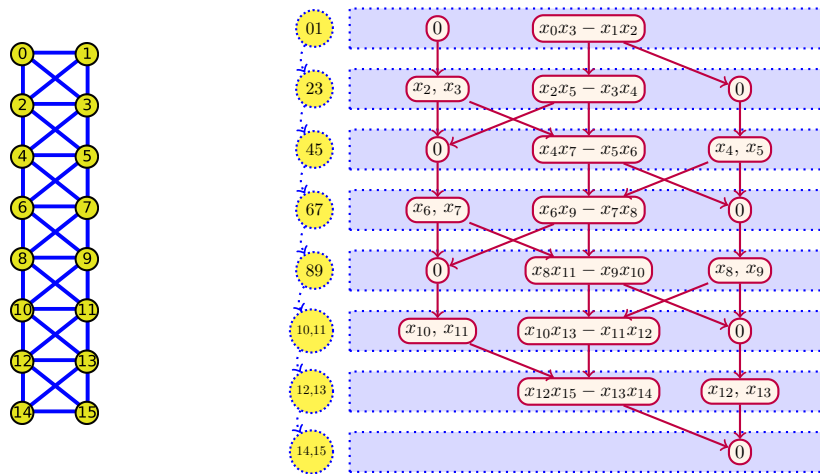


Figure 3. Chordal network for the ideal of adjacent minors.

As in Example 1.1, there is a simple recursive procedure to produce points on $\mathbf{V}(F_n)$: we choose the values of the last column of the matrix arbitrarily, and then for column i we either choose it arbitrarily, in the case that column $i + 1$ is zero, or scale column $i + 1$ if it is nonzero. This procedure is actually describing the irreducible components of the variety. In this way, the irreducible components admit a compact description, which is shown in Figure 3. Again, the components are given by the maximal directed paths (e.g., $T = \{0, x_2, x_3, 0, x_6, x_7, x_8x_{11} - x_9x_{10}, 0, x_{12}, x_{13}, 0\}$), and its cardinality is the n th Fibonacci number.

Contributions. The examples above show how certain polynomial systems with tree-like structure admit a compact chordal network representation. The aim of this paper is to develop a general framework to systematically understand and compute chordal networks. We also study how to effectively use chordal networks to solve different problems from computational algebraic geometry. A major difficulty is that exponentially many triangular sets may appear (e.g., the Fibonacci number in [Example 1.1](#)).

This paper presents the following contributions:

- We introduce the notion of chordal networks, a novel representation of polynomial ideals aimed toward exploiting structured sparsity.
- We develop the *chordal triangularization* method ([Algorithm 1](#)) to compute such chordal network representation. Its correctness is established in [Theorems 3.9](#) and [6.10](#).
- We show that several families of polynomial systems admit a “small” chordal network representation, of size $O(n)$. This is true for certain zero-dimensional ideals ([Remark 3.23](#)), all monomial ideals ([Theorem 5.3](#)), and certain binomial/determinantal ideals ([section 7.3](#)), although in general this cannot be guaranteed ([Remark 3.25](#)).
- We show how to effectively use chordal networks to compute several properties of the underlying variety. In particular, the cardinality ([section 4.2](#)), dimension, and top-dimensional component ([section 5.2](#)) can be computed in linear time. In some interesting cases we can also describe the irreducible components.
- We present a Monte Carlo algorithm to test radical ideal membership ([Algorithm 3](#)). We show in [Theorem 4.3](#) that the complexity is linear when the given polynomial preserves some of the graphical structure of the system.

We point out that we have a preliminary implementation of a Macaulay2 package with all the methods from this paper, and it is available at www.mit.edu/~diegcif.

Structure of the paper. The organization of this paper is as follows. In [section 2](#) we review the concept of chordal graph and then formalize the notion of chordal network. We then proceed to explain our methods, initially only for a restricted class of zero-dimensional problems ([section 3](#) and [section 4](#)), then for the case of monomial ideals ([section 5](#)), and finally considering the fully general case ([section 6](#)). We conclude the paper in [section 7](#) with numerical examples of our methods.

The reason for presenting our results in this stepwise manner is that the general case requires highly technical concepts from the theory of triangular sets. Indeed, we encourage the reader unfamiliar with triangular sets to omit [section 6](#) on the first read. On the other hand, by first specializing our methods to the zero-dimensional and monomial cases, we can introduce them all in a transparent manner. Importantly, the basic structure of the chordal triangularization algorithm, presented in [section 3](#), remains the same for the general case. Similarly, our algorithms that use chordal networks to compute properties of the variety (e.g., cardinality, dimension) introduced in [section 4](#) and [section 5](#) also extend in a natural way.

Related work. The development of chordal networks can be seen as a continuation of our earlier work [\[9\]](#), and we refer the reader to that paper for a detailed survey of the relevant literature on graphical structure in computational algebraic geometry. For this reason, below we only discuss related work in the context of triangular sets, and we point out the main differences between this paper and [\[9\]](#).

This paper improves upon [9] in two main areas. First, chordal networks provide a much richer description of the variety than the elimination ideals obtained by chordal elimination. For instance, the elimination ideals of the equations from Example 1.3 are trivial, but its chordal network representation reveals its irreducible components. In addition, neither the dimension, nor the cardinality, nor the radical ideal membership can be directly computed from the elimination ideals (we need a Gröbner basis). Second, we show how to compute chordal network representations for arbitrary polynomial systems (in characteristic zero). In contrast, chordal elimination only computes the elimination ideals under certain assumptions.

There is a broad literature studying triangular decompositions of ideals [3, 19, 24, 26, 36]. However, past work has not considered the case of sparse polynomial systems. Among the many existing triangular decomposition algorithms, Wang's elimination methods are particularly relevant to us [36, 37]. Although seemingly unnoticed by Wang, most of his algorithms preserve the chordal structure of the system. As a consequence, we have experimentally seen that his methods are more efficient than those based on regular chains [25, 26] for the examples considered in this paper.

As opposed to previous work, we emphasize chordal networks as our central object of study, rather than the explicit triangular decomposition obtained. This is a key distinction since for several families of ideals the size of the chordal network is linear even though the corresponding triangular decomposition has exponential size (see the examples from above). In addition, our methods deliberately treat triangular decompositions as a black box algorithm, allowing us to use either Lazard's LexTriangular algorithm [24] for the zero-dimensional case or Wang's RegSer algorithm [35] for the positive-dimensional case.

2. Chordal networks.

2.1. Chordal graphs. Chordal graphs have many equivalent characterizations. A good presentation is found in [4]. For our purposes, we use the following definition.

Definition 2.1. *Let G be a graph with vertices x_0, \dots, x_{n-1} . An ordering of its vertices $x_0 > x_1 > \dots > x_{n-1}$ is a perfect elimination ordering if for each x_l the set*

$$(2) \quad X_l := \{x_l\} \cup \{x_m : x_m \text{ is adjacent to } x_l, x_m < x_l\}$$

is such that the restriction $G|_{X_l}$ is a clique. A graph G is chordal if it has a perfect elimination ordering.

Remark 2.2. Observe that lower indices correspond to larger vertices.

Chordal graphs have many interesting properties. For instance, they have at most n maximal cliques, given that any clique is contained in some X_l . Note that trees are chordal graphs, since by successively pruning a leaf from the tree we get a perfect elimination ordering. We can always find a perfect elimination ordering of a chordal graph in linear time [28].

Definition 2.3. *Let \mathcal{G} be an arbitrary graph. We say that G is a chordal completion of \mathcal{G} if it is chordal and \mathcal{G} is a subgraph of G . The clique number of G , denoted as κ , is the size of its largest clique. The treewidth of \mathcal{G} is the minimum clique number of G (minus one) among all possible chordal completions.*

Observe that given any ordering $x_0 > \dots > x_{n-1}$ of the vertices of \mathcal{G} , there is a natural chordal completion G ; i.e., we add edges to \mathcal{G} in such a way that each $G|_{X_l}$ is a clique. In general, we want to find a chordal completion with a small clique number. However, there are $n!$ possible orderings of the vertices, and thus finding the best chordal completion is not simple. Indeed, this problem is NP-hard [2], but there are good heuristics and approximation algorithms [5, 32]. See [6] for a comparison of some of these heuristics.




Figure 4. Left: 10-vertex graph (blue/solid) and a chordal completion (green/dashed). Right: Elimination tree of the chordal completion.

Example 2.4. Let \mathcal{G} be the blue/solid graph in Figure 4a. This graph is not chordal, but if we add the six green/dashed edges shown in the figure, we obtain a chordal completion G . In fact, the ordering $x_0 > \dots > x_9$ is a perfect elimination ordering of the chordal completion. The clique number of G is four, and the treewidth of \mathcal{G} is three.

As mentioned earlier, we will assume throughout this document that the polynomial system F is supported on a given chordal graph G , where by supported we mean that G is a chordal completion of $\mathcal{G}(F)$. Moreover, we assume that the ordering of the vertices (inherited from the polynomial ring) is a perfect elimination ordering of G .

Given a chordal graph G with some perfect elimination ordering, there is an associated tree that will be very helpful in our discussion.

Definition 2.5. Let G be an ordered graph with vertices $x_0 > \dots > x_{n-1}$. The elimination tree of G is the following directed spanning tree: for each l there is an arc from x_l toward the largest x_p that is adjacent to x_l and $x_p < x_l$. We will say that x_p is the parent of x_l and x_l is a child of x_p . Note that the elimination tree is rooted at x_{n-1} .

Figure 4b shows an example of the elimination tree. We now present a simple property of the elimination tree of a chordal graph.

Lemma 2.6. Let G be a chordal graph, let x_l be some vertex, and let x_p be its parent in the elimination tree T . Then $X_l \setminus \{x_l\} \subseteq X_p$, where X_i is as in (2).

Proof. Let $Y := X_l \setminus \{x_l\}$. Note that Y is a clique, whose largest variable is x_p . Since X_p is the unique largest clique satisfying such a property, we must have $Y \subseteq X_p$. ■

2.2. Chordal networks. We proceed to formalize the concept of chordal networks.

Downloaded 06/28/19 to 18.30.8.120. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Definition 2.7. Let G be a chordal graph with vertex set X , and let X_l be as in (2). A G -chordal network is a directed graph \mathcal{N} , whose nodes are polynomial sets in $\mathbb{K}[X]$, such that the following hold:

- (Nodes supported on cliques) Each node F_l of \mathcal{N} is given a rank $l = \text{rk}(F_l)$, with $0 \leq l < n$, such that $F_l \subseteq \mathbb{K}[X_l]$.
- (Arcs follow elimination tree) If (F_l, F_p) is an arc of \mathcal{N} , then (l, p) is an arc of the elimination tree of G , where $l = \text{rk}(F_l), p = \text{rk}(F_p)$.

A chordal network is triangular if each node consists of a single polynomial f , and either $f = 0$ or its largest variable is $x_{\text{rk}(f)}$.

There is one parameter of a chordal network that will determine the complexity of some of our methods. The *width* of a chordal network, denoted as W , is the largest number of nodes of any given rank. Note that the number of nodes in the network is at most nW , and the number of arcs is at most $(n - 1)W^2$.

We can represent chordal networks using the diagrams we have shown throughout the paper. Since the structure of a chordal network resembles the elimination tree (second item in the definition), we usually show the elimination tree to the left of the network.

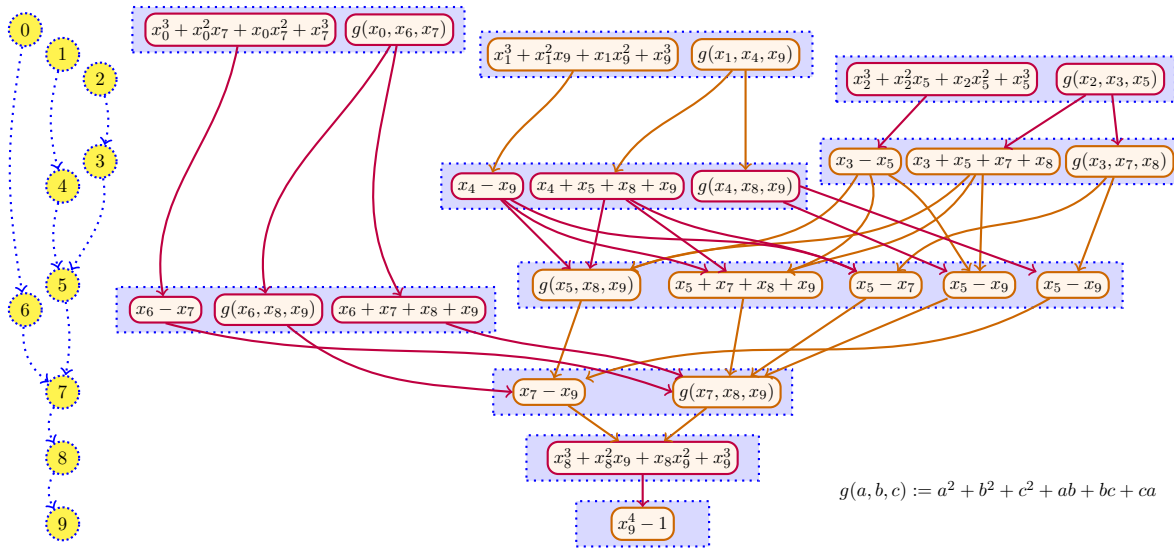


Figure 5. G -chordal network, where G is the chordal graph from Figure 4a. The elimination tree of G is shown on the left.

Example 2.8. Let \mathcal{G} be the blue/solid graph from Figure 4a, and let G be the green/dashed chordal completion. Figure 5 shows a G -chordal network of width 5 that represents the 4-colorings of graph \mathcal{G} (equation (1)). The elimination tree of G is shown to the left of the diagram. Note that this network is triangular, and thus all its nodes consist of a single polynomial. For instance, two of its nodes are $f_5 = x_5 + x_7 + x_8 + x_9$ and $f_6 = x_6 - x_7$. Nodes are grouped in blue rectangular boxes according to their rank. In particular, f_5 has rank 5 and f_6 rank 6, and indeed $f_5 \in \mathbb{K}[X_5] = \mathbb{K}[x_5, x_7, x_8, x_9]$ and $f_6 \in \mathbb{K}[X_6] = \mathbb{K}[x_6, x_7, x_8, x_9]$.

Example 2.9. Let \mathcal{G} be the 9-cycle with vertices x_0, \dots, x_8 . Let G be the chordal completion obtained by connecting vertex x_8 to all the others. [Figure 1](#) shows a triangular G -chordal network. The elimination tree, shown to the left of the network, is the path $x_0 \rightarrow \dots \rightarrow x_8$.

Remark 2.10. Sometimes we collapse certain ranks to make the diagram visually simpler. In particular, in [Figure 3](#) we collapse the ranks $2i, 2i + 1$ into a single group.

As suggested by the examples in the introduction, a triangular chordal network gives a decomposition of the polynomial ideal into triangular sets. Each such triangular set corresponds to a *chain* of the network, as defined next.

Definition 2.11. Let \mathcal{N} be a G -chordal network. A chain of \mathcal{N} is a tuple of nodes $C = (F_0, F_1, \dots, F_{n-1})$ such that the following hold:

- $\text{rk}(F_l) = l$ for each l .
- If x_p is the parent of x_l , then (F_l, F_p) is an arc of \mathcal{N} .

Example 2.12. The chordal network from [Figure 5](#) has 21 chains, one of which is

$$C = (x_0^2 + x_0x_6 + x_0x_7 + x_6^2 + x_6x_7 + x_7^2, x_1^3 + x_1^2x_9 + x_1x_9^2 + x_9^3, x_2^3 + x_2^2x_5 + x_2x_5^2 + x_5^3, \\ x_3 - x_5, x_4 - x_9, x_5^2 + x_5x_8 + x_5x_9 + x_8^2 + x_8x_9 + x_9^2, \\ x_6^2 + x_6x_8 + x_6x_9 + x_8^2 + x_8x_9 + x_9^2, x_7 - x_9, x_8^3 + x_8^2x_9 + x_8x_9^2 + x_9^3, x_9^4 - 1).$$

2.3. Binary decision diagrams. Although motivated from a different perspective and with quite distinct goals, throughout the development of this paper we realized the intriguing similarities between chordal networks and a class of data structures used in computer science known as ordered binary decision diagrams (OBDDs) [[1, 7, 21, 38](#)].

A binary decision diagram (BDD) is a data structure that can be used to represent Boolean (binary) functions in terms of a directed acyclic graph. They can be interpreted as a binary analogue of a straight-line program, where the nodes are associated with variables and the outgoing edges of a node correspond to the possible values of that variable. A particularly important subclass are the OBDDs, where the branching occurs according to a specific fixed variable ordering. Under a mild condition (reducibility) this representation can be made unique, and thus every Boolean function has a canonical OBDD representation. OBDDs can be effectively used for further manipulation (e.g., decide satisfiability, count satisfying assignments, compute logical operations). Interestingly, several important functions have a compact OBDD representation. A further variation, *zero-suppressed* BDDs (ZBDDs), can be used to efficiently represent subsets of the hypercube $\{0, 1\}^n$ and to manipulate them (e.g., intersections, sampling, linear optimization).

Chordal networks can be thought of as a wide generalization of OBDDs/ZBDDs to arbitrary algebraic varieties over general fields (instead of finite sets in $(\mathbb{F}_2)^n$). Like chordal networks, an OBDD corresponds to a certain directed graph, but where the nodes are variables (x_0, x_1, \dots) instead of polynomial sets. We will see in [section 5](#) that for the specific case of monomial ideals, the associated chordal networks also have this form. Since one of our main goals is to preserve graphical structure for efficient computation, in this paper we define chordal networks only for systems that are structured according to some chordal graph. In addition, for computational purposes we do not insist on uniqueness of the representation (although it might be possible to make them canonical after further processing).

The practical impact of data structures like BDDs and OBDDs over the past three decades has been very significant, as they have enabled breakthrough results in many areas of computer science, including model checking, formal verification, and logic synthesis. We hope that chordal networks will make possible similar advances in computational algebraic geometry. The connections between BDDs and chordal networks run much deeper, and we plan to further explore them in the future.

3. The chordally zero-dimensional case. In this section we present our main methods to compute triangular chordal networks, although focused on a restricted type of zero-dimensional problems. This restriction is for simplicity only; we will see that our methods naturally extend to arbitrary ideals. Concretely, we consider the following family of polynomial sets.

Definition 3.1 (chordally zero-dimensional). *Let $F \subseteq \mathbb{K}[X]$ be supported on a chordal graph G . We say that F is chordally zero-dimensional if for each maximal clique X_l of graph G the ideal $\langle F \cap \mathbb{K}[X_l] \rangle$ is zero-dimensional.*

Note that the q -coloring equations in (1) are chordally zero-dimensional. As in Example 1.1, chordally zero-dimensional problems always have simple chordal network representations.

Remark 3.2 (the geometric picture). There is a nice geometric interpretation behind the chordally zero-dimensional condition. Denoting V_l the variety of $F \cap \mathbb{K}[X_l]$, the condition is that each V_l is finite. Note now that $\pi_{X_l}(\mathbf{V}(F)) \subseteq V_l$, where π_{X_l} denotes the projection onto the coordinates of X_l . Thus, independent of the size of $\mathbf{V}(F)$, the chordally zero-dimensional condition allows us to bound the size of its projections onto each X_l . More generally, although not elaborated in this paper, our methods are expected to perform well on any F (possibly positive-dimensional) for which the projections $\pi_{X_l}(\mathbf{V}(F))$ are well behaved.

3.1. Triangular sets. We now recall the basic concepts of triangular sets for the case of zero-dimensional ideals, following [24]. We delay the exposition of the positive-dimensional case to section 6.

Definition 3.3. *Let $f \in \mathbb{K}[X] \setminus \mathbb{K}$ be a nonconstant polynomial. The main variable of f , denoted $\text{mvar}(f)$, is the greatest variable appearing in f . The initial of f , denoted $\text{init}(f)$, is the leading coefficient of f when viewed as a univariate polynomial in $\text{mvar}(f)$. A zero-dimensional triangular set is a collection of nonconstant polynomials $T = \{t_0, \dots, t_{n-1}\}$ such that $\text{mvar}(t_i) = x_i$ and $\text{init}(t_i) = 1$ for each i .*

Most of the analysis done in this paper will work over an arbitrary field \mathbb{K} . For some results we require the field to contain sufficiently many elements, so we might need to consider a field extension. We denote by $\overline{\mathbb{K}}$ the algebraic closure of \mathbb{K} . For a polynomial set F , we let $\mathbf{V}(F) \subseteq \overline{\mathbb{K}}^n$ be its variety. Note that for a zero-dimensional triangular set T , we always have

$$(3) \quad |\mathbf{V}(T)| \leq \deg(T) := \prod_{t \in T} \text{mdeg}(t),$$

where $\text{mdeg}(t) := \deg(t, \text{mvar}(t))$ denotes the degree on the main variable. Furthermore, the above is an equality if we count multiplicities.

For a triangular set T , let $\langle T \rangle$ denote the generated ideal. It is easy to see that a zero-dimensional triangular set is a lexicographic Gröbner basis of $\langle T \rangle$. In particular, we can test ideal membership by taking normal form. We also denote as $\text{elim}_p(T) := T \cap \mathbb{K}[x_p, x_{p+1}, \dots]$ the subset of T restricted to variables less than or equal to x_p . Note that $\text{elim}_p(T)$ generates the elimination ideal of $\langle T \rangle$ because of the elimination property of lexicographic Gröbner bases.

Notation. We let $S = S_1 \sqcup S_2$ denote a disjoint union, i.e., $S = S_1 \cup S_2$ and $S_1 \cap S_2 = \emptyset$.

Definition 3.4. Let $I \subseteq \mathbb{K}[X]$ be a zero-dimensional ideal. A triangular decomposition of I is a collection \mathcal{T} of triangular sets, such that $\mathbf{V}(I) = \bigsqcup_{T \in \mathcal{T}} \mathbf{V}(T)$. We say that \mathcal{T} is square-free if each $T \in \mathcal{T}$ generates a radical ideal. We say that \mathcal{T} is irreducible if each $T \in \mathcal{T}$ generates a prime ideal (or, equivalently, a maximal ideal).

Lazard proposed algorithms to compute a triangular decomposition from a Gröbner basis [24]. He also showed how to postprocess it to make it square-free/irreducible.

Remark 3.5. As explained in [24], there might be several distinct triangular decompositions of an ideal, but there are simple ways to pass from one description to another.

3.2. Chordal triangularization. We proceed to explain how to compute a triangular chordal network representation of a polynomial set F . We will start with a particular (induced) chordal network that is modified step after step to make it triangular.

Definition 3.6. Let $F \subseteq \mathbb{K}[X]$ be supported on a chordal graph G . The induced G -chordal network has a unique node of rank k , namely $F_k := F \cap \mathbb{K}[X_k]$, and its arcs are the same as in the elimination tree; i.e., (F_l, F_p) is an arc if x_p is the parent of x_l .

We will sequentially perform two types of operations to the induced chordal network.

Triangulate(F_l). Let \mathcal{T} be a triangular decomposition of a node F_l of the network. Replace node F_l with one node for each triangular set in \mathcal{T} . Any node which was previously connected to F_l is then connected to each of the new nodes.

Eliminate(T). Let T be a rank l node, and let x_p be the parent of x_l . Let $T_p := \text{elim}_p(T)$ and $T_l := T \setminus T_p$. For each arc (T, F_p) we create a new rank p node $F'_p := F_p \cup T_p$, and we substitute arc (T, F_p) with (T, F'_p) . Then, we copy all arcs coming out of F_p to F'_p (while keeping the old arcs). Next, we replace the content of node T with the polynomial set T_l .

The operations are performed in rounds: in the l th round we triangulate/eliminate all rank l nodes. After each round, we may reduce the network with the following additional operations.

MergeIn(l). Merge any two rank l nodes F_l, F'_l if they define the same ideal, and they have the same sets of incoming arcs.

MergeOut(l). Merge any two rank l nodes F_l, F'_l if they define the same ideal, and they have the same sets of outgoing arcs.

Example 3.7. Consider the polynomial set $F = \{x_0^3 - x_0, x_0x_2 - x_2, x_1 - x_2, x_2^2 - x_2, x_2x_3^2 - x_3\}$, whose associated graph is the star graph (x_2 is connected to x_0, x_1, x_3). Figure 6 illustrates a sequence of operations (triangulation, elimination, and merge) performed on its induced

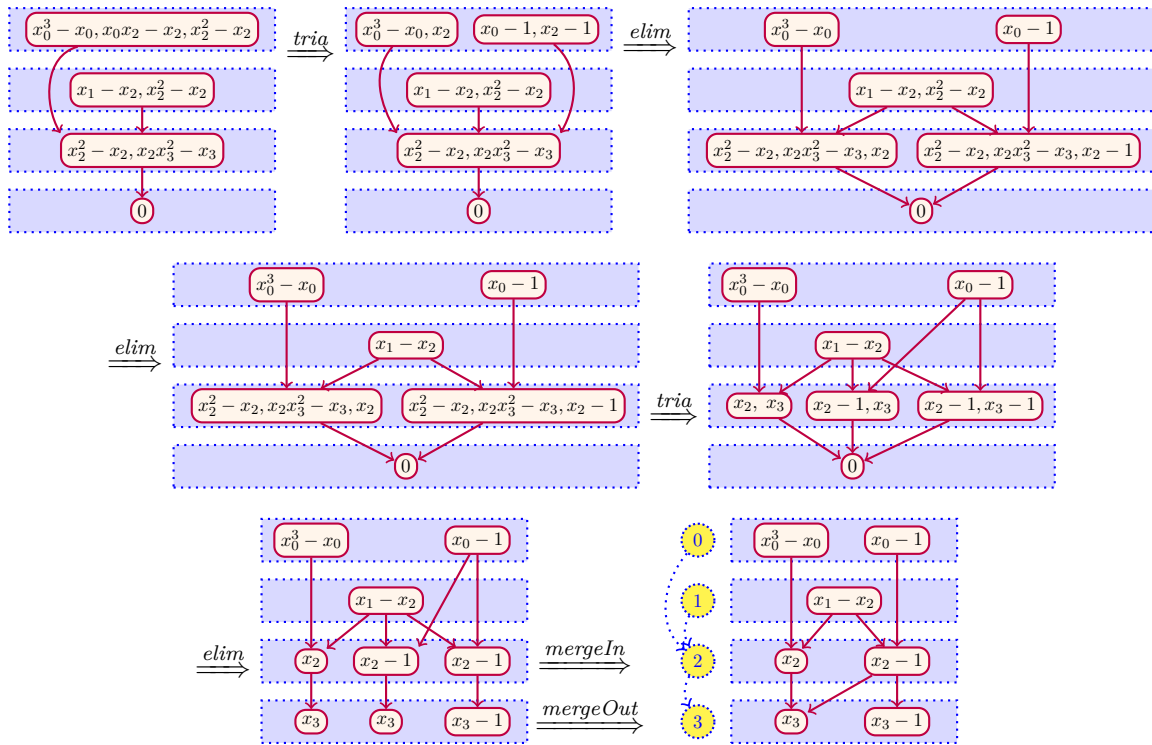


Figure 6. Chordal triangularization from Example 3.7.

chordal network. The chordal network obtained has three chains:

$$(x_0^3 - x_0, x_1 - x_2, x_2, x_3), (x_0 - 1, x_1 - x_2, x_2 - 1, x_3), (x_0 - 1, x_1 - x_2, x_2 - 1, x_3 - 1).$$

These chains give a triangular decomposition of F .

Algorithm 1 presents the chordal triangularization method. The input consists of a polynomial set $F \subseteq \mathbb{K}[X]$ and a chordal graph G . As in the above example, the output of the algorithm is always a triangular chordal network, and it encodes a triangular decomposition of the given polynomial set F .

3.3. Algorithm analysis. The objective of this section is to prove that, when the input F is chordally zero-dimensional (Definition 3.1), Algorithm 1 produces a G -chordal network, whose chains give a triangular decomposition of F . As described below, the chordally zero-dimensional assumption is only needed in order for the algorithm to be *well defined* (recall that up to this point we have only defined triangular decompositions of zero-dimensional systems). Later in the paper we will see how to extend Algorithm 1 to arbitrary ideals.

Definition 3.8. Let \mathcal{N} be a chordal network, and let $C = (F_0, \dots, F_{n-1})$ be a chain. The variety of the chain is $\mathbf{V}(C) := \mathbf{V}(F_0 \cup \dots \cup F_{n-1})$. The variety $\mathbf{V}(\mathcal{N})$ of the chordal network is the union of $\mathbf{V}(C)$ over all chains C .

Theorem 3.9. Let $F \subseteq \mathbb{K}[X]$, supported on chordal graph G , be chordally zero-dimensional. Algorithm 1 computes a G -chordal network whose chains give a triangular decomposition of F .

Algorithm 1 Chordal triangularization**Input:** Polynomial set $F \subseteq \mathbb{K}[X]$ supported on a chordal graph G **Output:** Triangular G -chordal network \mathcal{N} such that $\mathbf{V}(\mathcal{N}) = \mathbf{V}(F)$

```

1: procedure CHORDALNET( $F, G$ )
2:    $\mathcal{N} :=$  induced  $G$ -chordal network of  $F$ 
3:   for  $l = 0 : n - 1$  do
4:     for  $F_l$  node of  $\mathcal{N}$  of rank  $l$  do
5:       TRIANGULATE( $F_l$ )
6:       MERGEOUT( $l$ )
7:       if  $l < n - 1$  then
8:          $x_p :=$  parent of  $x_l$ 
9:         for  $T$  node of  $\mathcal{N}$  of rank  $l$  do
10:          ELIMINATE( $T$ )
11:          MERGEOUT( $p$ )
12:        MERGEIN( $l$ )
13:   return  $\mathcal{N}$ 

```

We will split the proof of [Theorem 3.9](#) into several lemmas. We first show that the algorithm is well defined; i.e., we only perform triangulation operations (line 5) on nodes F_l that define zero-dimensional ideals.

Lemma 3.10. *Let $F \subseteq \mathbb{K}[X]$ be chordally zero-dimensional. Then in [Algorithm 1](#) every triangulation operation is performed on a zero-dimensional ideal.*

Proof. See [Appendix A.1](#). ■

We now show that the chordal structure is preserved during the algorithm.

Lemma 3.11. *Let \mathcal{N} be a G -chordal network. Then the result of performing a triangulation or elimination operation is also a G -chordal network.*

Proof. Consider first a triangulation operation. Note that if $F_l \subseteq \mathbb{K}[X_l]$, then each T in a triangular decomposition is also in $\mathbb{K}[X_l]$. Consider now an elimination operation. Let $T \subseteq \mathbb{K}[X_l]$ and $F_p \subseteq \mathbb{K}[X_p]$ be two adjacent nodes. Using [Lemma 2.6](#), $\text{elim}_p(T) \subseteq \mathbb{K}[X_l \setminus \{x_l\}] \subseteq \mathbb{K}[X_p]$. Thus, the new node $F'_p := F_p \cup \text{elim}_p(T) \subseteq \mathbb{K}[X_p]$. It is clear that for both operations the layered structure of \mathcal{N} is preserved (i.e., arcs follow the elimination tree). ■

We next show that the chains of the output network are triangular sets.

Lemma 3.12. *The output of [Algorithm 1](#) is a triangular G -chordal network.*

Proof. Let $T \subseteq \mathbb{K}[X_l]$ be a rank l node for which we will perform an elimination operation. Note that T must be triangular as we previously performed a triangulation operation. Therefore, there is a unique polynomial $f \in T$ with $\text{mvar}(f) = x_l$. When we perform the elimination operation this is the only polynomial of T we keep, which concludes the proof. ■

Finally, we show that the variety is preserved during the algorithm.

Lemma 3.13. *Let \mathcal{N} be the output of [Algorithm 1](#). Then $\mathbf{V}(\mathcal{N}) = \mathbf{V}(F)$, and moreover, any two chains of \mathcal{N} have disjoint varieties.*

Proof. Let us show that the variety is preserved when we perform triangulation, elimination, and merge operations. First, note that a merge operation does not change the set of chains of the network, so the variety is preserved. Consider now the case of a triangulation operation. Let \mathcal{N} be a chordal network, and let F be one of its nodes. Let \mathcal{T} be a triangular decomposition of F , and let \mathcal{N}' be the chordal network obtained after replacing F with \mathcal{T} . Let C be a chain of \mathcal{N} containing F , and let $C' = C \setminus \{F\}$. Then

$$\mathbf{V}(C) = \mathbf{V}(C') \cap \mathbf{V}(F) = \mathbf{V}(C') \cap \left(\bigsqcup_{T \in \mathcal{T}} \mathbf{V}(T) \right) = \bigsqcup_{T \in \mathcal{T}} \mathbf{V}(C') \cap \mathbf{V}(T) = \bigsqcup_{T \in \mathcal{T}} \mathbf{V}(C' \cup \{T\}).$$

Note that $C' \cup \{T\}$ is a chain of \mathcal{N}' . Moreover, all chains of \mathcal{N}' that contain one of the nodes of \mathcal{T} have this form. Thus, the triangulation step indeed preserves the variety.

Finally, consider the case of an elimination operation. Let $T \subseteq \mathbb{K}[X_l]$ be a node, let (T, F_p) be an arc, and let $F'_p = F_p \cup \text{elim}_p(T)$, $T_l = T \setminus \text{elim}_p(T)$. Let \mathcal{N}' be the network obtained after an elimination step on T . It is clear that

$$\mathbf{V}(T \cup F_p) = \mathbf{V}(T_l \cup \text{elim}_p(T) \cup F_p) = \mathbf{V}(T_l \cup F'_p).$$

Since a chain in \mathcal{N} containing T, F_p turns into a chain in \mathcal{N}' containing T_l, F'_p , we conclude that the elimination step also preserves the variety. ■

Proof of [Theorem 3.9](#). We already proved the theorem, since we showed that the algorithm is well defined ([Lemma 3.10](#)), chordal structure is preserved ([Lemma 3.11](#)), and the chains in the output are triangular sets ([Lemma 3.12](#)) that decompose the given variety ([Lemma 3.13](#)). ■

3.4. Radical and irreducible decompositions. We just showed that [Algorithm 1](#) can compute chordal network representations of some zero-dimensional problems. However, we sometimes require additional properties of the chordal network. In particular, in [section 4](#) we will need square-free representations, i.e., such that any chain generates a radical ideal. As shown next, we can obtain such representations by making one change in [Algorithm 1](#): whenever we perform a triangulation operation, we should produce a square-free decomposition.

Proposition 3.14. *Assume that all triangular decompositions computed in [Algorithm 1](#) are square-free. Then any chain of the output network generates a radical ideal.*

Proof. See [Appendix A.1](#). ■

Instead of radicality, we could further ask for an irreducible representation, i.e., such that any chain generates a prime ideal. The obvious modification to make is to require all triangulation operations to produce irreducible decompositions. Unfortunately, this does not always work. Indeed, we can find irreducible univariate polynomials $f \subseteq \mathbb{K}[x_0]$, $g \subseteq \mathbb{K}[x_1]$ such that $\langle f, g \rangle \subseteq \mathbb{K}[x_0, x_1]$ is not prime (e.g., $f = x_0^2 + 1$, $g = x_1^2 + 1$).

Nonetheless, there is an advantage of maintaining prime ideals through the algorithm: it gives a simple bound on the size of the triangular network computed, as shown next. This bound will be used when analyzing the complexity of the algorithm.

Lemma 3.15. *Assume that all triangular decompositions computed in Algorithm 1 are irreducible. Then the number of rank l nodes in the output is at most $|\mathbf{V}(F \cap \mathbb{K}[X_l])|$.*

Proof. Let us see that there are at most $|\mathbf{V}(F \cap \mathbb{K}[X_l])|$ rank l nodes after the merge operation from line 6. First note that when we perform this operation any rank l node has an outgoing arc to all rank p nodes (where x_p is the parent of x_l). Therefore, this operation merges any two rank l nodes that define the same ideal. Since these ideals are all maximal, then for any two distinct nodes T_l, T'_l we must have $\mathbf{V}(T_l) \cap \mathbf{V}(T'_l) = \emptyset$. Also note that both $\mathbf{V}(T_l), \mathbf{V}(T'_l)$ are subsets of $\mathbf{V}(F \cap \mathbb{K}[X_l])$. The lemma follows. ■

Remark 3.16. There are other ways to achieve the above bound that do not require computing irreducible decompositions. For instance, we can force the varieties $\mathbf{V}(T_l), \mathbf{V}(T'_l)$ to be disjoint by using ideal saturation.

3.5. Complexity. We proceed to estimate the cost of Algorithm 1 in the chordally zero-dimensional case. We will show that the complexity² is $O(nq^{O(\kappa)})$, where κ is the treewidth (or clique number) of the graph, and q is a certain degree bound on the polynomials that we formalize below. In particular, when the treewidth κ is bounded, the complexity is linear in n and polynomial in the degree bound q .

Definition 3.17 (q -domination). *We say that a polynomial set $F_l \subseteq \mathbb{K}[X_l]$ is q -dominated if for each $x_i \in X_l$ there is some $f \in F_l$ such that $\text{mvar}(f) = x_i$, $\text{init}(f) = 1$, and $\text{deg}(f, x_i) \leq q$. Let $F \subseteq \mathbb{K}[X]$ be supported on a chordal graph G . We say that F is chordally q -dominated if $F \cap \mathbb{K}[X_l]$ is q -dominated for each maximal clique X_l of graph G .*

Example 3.18. The coloring equations in (1) are chordally q -dominated since the equations $x_i^q - 1$ are present. Another important example is the case of finite fields \mathbb{F}_q , since if we include the equations $x_i^q - x_i$, as is often done, the problem becomes chordally q -dominated.

Remark 3.19. Observe that if F is chordally q -dominated, then it is also chordally zero-dimensional. Conversely, if F is chordally zero-dimensional, then we can apply a simple transformation to make it chordally q -dominated (for some q). Concretely, for each maximal clique X_l we can enlarge F with a Gröbner basis of $F \cap \mathbb{K}[X_l]$.

We note that we also used the q -dominated condition in [9] to analyze the complexity of chordal elimination. The importance of this condition is that it allows us to easily bound the complexity of computing Gröbner bases or triangular decompositions, as stated next.

Proposition 3.20. *For any q -dominated polynomial set on k variables, the complexity of computing Gröbner bases and (square-free/irreducible) triangular decompositions is $q^{O(k)}$.*

Proof. See Appendix A.1. ■

The above proposition gives us the cost of the triangulation operations. However, we need to ensure that these operations are indeed performed on a q -dominated ideal, as shown next.

Lemma 3.21. *Let $F \subseteq \mathbb{K}[X]$ be chordally q -dominated. Then in Algorithm 1 any triangulation operation is performed on a q -dominated ideal.*

Proof. The proof is analogous to that of Lemma 3.10. ■

²Here the complexity is measured in terms of the number of field operations.

We are ready to estimate the complexity of chordal triangularization. For the analysis we assume that the merge operation from line 6 (resp., line 11) is performed simultaneously with the triangulation (resp., elimination) operations; i.e., as soon as we create a new node, we compare it with the previous nodes of the same rank to check if it is repeated.

Lemma 3.22. *Let $F \subseteq \mathbb{K}[X]$ be chordally q -dominated. Assume that all triangular decompositions computed in Algorithm 1 are irreducible. Then throughout the algorithm the width of the network is always bounded by q^κ , independently of the number of variables.*

Proof. This is a consequence of Lemma 3.15. See Appendix A.1. ■

Remark 3.23 (chordal network of linear size). It follows from the lemma that for fixed q, κ , any chordally q -dominated $F \subseteq \mathbb{K}[X]$ of treewidth κ has a chordal network representation with $O(n)$ nodes.

Theorem 3.24. *Let $F \subseteq \mathbb{K}[X]$ be chordally q -dominated. The complexity of chordal triangularization is $O(nWq^{O(\kappa)})$, where W is a bound on the width of the network throughout the algorithm. If all triangulation operations are irreducible, the complexity is $O(nq^{O(\kappa)})$.*

Proof. From Proposition 3.20 and Lemma 3.21 we know that each triangulation operation takes $q^{O(\kappa)}$, and thus the cost of all triangulations is $O(nWq^{O(\kappa)})$. The cost of the elimination operations is negligible. As for the merging operations, we can efficiently verify if a new node is repeated by using a hash table. Thus, the cost of the merging operation is also negligible. Finally, if all triangulation operations are irreducible, then $W \leq q^k$ because of Lemma 3.22. ■

Remark 3.25 (beyond chordally zero-dimensional). We will later see that, after a suitable redefinition of the triangulation step, Algorithm 1 can also be applied to arbitrary ideals. Nonetheless, the complexity bounds from above do depend on the special structure of the chordally zero-dimensional case. Indeed, solving polynomial equations of treewidth one is NP-hard [9, Ex. 1.1], and counting their number of solutions is #P-hard (even in the generic case for treewidth two [8, Prop. 24]). As a consequence, chordal triangularization will not always run in polynomial time. When using Algorithm 1 in such hard instances, we may end up with very high degree polynomials or with a very large number of nodes.

4. Computing with chordal networks. Triangular decompositions are one of the most common tools in computational algebraic geometry. The reason is that there are many good algorithms to compute them and that they can be used to derive several properties of the underlying variety. However, as seen in Example 1.1, the size of the decomposition obtained might be extremely large (exponential) even for very simple cases. Chordal networks can provide a compact representation for these large decompositions. We will see how to effectively use the data structure of chordal networks to compute several properties of the variety.

Let $I = \langle F \rangle$ be a zero-dimensional ideal. We consider the following problems.

Elimination. Describe the projection of $\mathbf{V}(I)$ onto the last $n - l$ coordinates.

Zero count. Determine the number of solutions, i.e., the cardinality of $\mathbf{V}(I)$.

Sampling. Sample random points from $\mathbf{V}(I)$ uniformly.

Radical membership. Determine if a polynomial $h \in \mathbb{K}[X]$ vanishes on $\mathbf{V}(I)$, or, equivalently, determine if $h \in \sqrt{I}$.

In this section we will develop efficient algorithms for the above problems, given a square-

free chordal network \mathcal{N} (with possibly exponentially many chains). Recall that such a network can be obtained as explained in Proposition 3.14. We will see that the first three problems can be solved relatively easily. The radical membership problem is more complicated, and most of this section will be dedicated to it. We note that the algorithms for elimination and radical membership will naturally extend to the positive-dimensional case.

4.1. Elimination. The elimination problem is particularly simple, thanks to the elimination property of lexicographic Gröbner bases. For an arbitrary chordal network \mathcal{N} , let $\mathcal{N}_{\geq l}$ denote the subset of \mathcal{N} consisting of nodes of rank k with $k \geq l$. Then $\mathcal{N}_{\geq l}$ is a chordal network representation of the projection of $\mathbf{V}(I)$ onto the last $n - l$ coordinates.

4.2. Counting solutions. We want to determine $|\mathbf{V}(\mathcal{N})|$ for a square-free chordal network \mathcal{N} . Recall from (3) that $|\mathbf{V}(T)| = \deg(T)$ for a square-free triangular set T . Therefore, we just need to compute the sum of $\deg(C)$ over all chains C of the network. We can do this efficiently via dynamic programming, as explained in the following example.

Example 4.1 (zero count). Let us determine $|\mathbf{V}(\mathcal{N})|$ for the chordal network from Figure 5, which corresponds to counting 4-colorings for the blue/solid graph from Figure 4a. For a rank l node f_l of the network, let its weight $w(f_l)$ be its degree in x_l . Then we just need to compute $\sum_C \prod_{f_l \in C} w(f_l)$ where the sum is over all chains of the network. We can do this efficiently by successively eliminating the nodes of the network.

Let us first eliminate the nodes of rank 0. Let f_0^a, f_0^b be the two nodes of rank 0, with weights $w(f_0^a) = 3, w(f_0^b) = 2$. Let f_6^a, f_6^b, f_6^c be the nodes of rank 6, with weights $w(f_6^a) = w(f_6^c) = 1, w(f_6^b) = 2$. Note that any chain containing f_6^a must also contain f_0^a . Therefore, we can remove the arc (f_0^a, f_6^a) and update the weight $w(f_6^a) = 1 \times 3$. Similarly, any chain containing f_6^b (or f_6^c) must also contain f_0^b . So we may delete the arcs (f_0^b, f_6^b) and (f_0^b, f_6^c) and update the weights $w(f_6^b) = 2 \times 2, w(f_6^c) = 1 \times 2$. By doing this, we have disconnected, or eliminated, all nodes of rank 0. Continuing this procedure, the final weights obtained for each rank are shown below. The number of solutions is the last number computed: 10968.

$$\begin{aligned} \text{rk}(0) &\rightarrow [3, 2], & \text{rk}(1) &\rightarrow [3, 2], & \text{rk}(2) &\rightarrow [3, 2], & \text{rk}(3) &\rightarrow [3, 2, 4], \\ \text{rk}(4) &\rightarrow [3, 2, 4], & \text{rk}(5) &\rightarrow [50, 25, 20, 20, 16], & \text{rk}(6) &\rightarrow [3, 4, 2], \\ \text{rk}(7) &\rightarrow [264, 650], & \text{rk}(8) &\rightarrow [2742], & \text{rk}(9) &\rightarrow [10968]. \end{aligned}$$

Algorithm 2 generalizes the above example to arbitrary chordal networks. The complexity is $O(nW^2)$, since we perform one operation for each arc of the network.

4.3. Sampling solutions. Uniformly sampling solutions can be done quite easily by using the partial root counts computed in Algorithm 2. Instead of giving a formal description we simply illustrate the procedure with an example.

Example 4.2 (sampling). Consider again the chordal network of Figure 5. We want to uniformly sample a point $(\hat{x}_0, \dots, \hat{x}_9)$ from its variety, and we follow a bottom-up strategy. Let us first choose the value \hat{x}_9 . Since there is a unique rank 9 node $f_9 = x_9^4 - 1$, then \hat{x}_9 must be one of its four roots. Note that each of these roots extends to 2742 solutions (a fourth of the total number of solutions). Therefore, \hat{x}_9 should be equally likely to be any of these roots. Given the value of \hat{x}_9 , we can now set \hat{x}_8 to be any of the three roots of

Algorithm 2 Count solutions**Input:** Chordal network \mathcal{N} (triangular, squarefree)**Output:** Cardinality of $\mathbf{V}(\mathcal{N})$

```

1: procedure ZEROCOUNT( $\mathcal{N}$ )
2:   for  $f$  node of  $\mathcal{N}$  do
3:      $w(f) := (\text{mdeg}(f)$  if  $x_{\text{rk}(f)}$  is a leaf else 0)
4:   for  $l = 0 : n - 1$  do
5:     for  $(f_l, f_p)$  arc of  $\mathcal{N}$  with  $\text{rk}(f_l) = l$  do
6:        $w(f_p) := w(f_p) + w(f_l) \text{mdeg}(f_p)$ 
7:   return sum of  $w(f_{n-1})$  over all nodes of rank  $n - 1$ 

```

$f_8 = x_8^3 + x_8^2 x_9 + x_8 x_9^2 + x_9^3$, each equally likely. Consider now the two rank 7 nodes f_7^a, f_7^b of degrees 1 and 2. Note that \hat{x}_7 should be either a root of f_7^a or a root of f_7^b (for the given values of \hat{x}_8, \hat{x}_9). In order to sample uniformly, we need to know the number of solutions that each of those values extends to. From [Example 4.1](#) we know that f_7^a leads to 264 points on the variety, and f_7^b leads to 650. Therefore, we can decide which of them to use based on those weights. Assuming we choose f_7^b , we can now set \hat{x}_7 to be any of its two roots, each equally likely. It is clear how to continue.

4.4. Radical membership. In the radical ideal membership problem we want to check whether $h \in \mathbb{K}[X]$ vanishes on $\mathbf{V}(\mathcal{N})$. This is equivalent to determining whether for each chain C of \mathcal{N} the normal form $h_C := h \bmod C$ is identically zero. We will propose a Monte Carlo algorithm to efficiently test this property (without iterating over all chains) under certain structural assumptions on the polynomial h . Our main result is the following.

Theorem 4.3 (radical membership). *Let $F \subseteq \mathbb{K}[X]$ be chordally q -dominated. Let \mathcal{N} be a chordal network representation of F of width W . Let h be a polynomial that decomposes as $h = \sum_l h_l$ with $h_l \in \mathbb{K}[X_l]$. There is a Monte Carlo algorithm that determines whether h vanishes on $\mathbf{V}(F)$ in $\tilde{O}(nWq^{2\kappa} + nW^2q^\kappa)$. Here the notation \tilde{O} ignores polynomial factors in the clique number κ .*

Remark 4.4. The theorem is restricted to polynomials h that preserve some of the structure of the graph G , although they may involve all the variables in the ring $\mathbb{K}[X]$ (as opposed to the polynomials of the chordal network). The above-mentioned Monte Carlo algorithm also works for other types of polynomials h , but we do not prove complexity bounds for them.

We point out that the above complexity result is far from trivial. To justify this claim we can show that a simple variation of the radical membership problem is NP-hard under very mild assumptions.

Example 4.5 (zero divisor problem). Consider the zero divisor problem: determine if a polynomial $h \in \mathbb{K}[X]$ vanishes on at least one point of $\mathbf{V}(I)$. Also consider the NP-complete subset sum problem: decide if a set of integers $A = \{a_0, \dots, a_{n-1}\}$ contains a subset whose sum is some given value S . We can reduce it to the zero divisor problem by considering the ideal $I := \langle x_i(x_i - a_i) : 0 \leq i < n \rangle$ and the polynomial $h := \sum_i x_i - S$. Note that the

associated graph is the completely disconnected graph ($\kappa = 1$), and thus its induced chordal network is already triangular ($W = 1, q = 2$).

We proceed to derive our radical ideal membership test. We will initially assume that the variables of h are all contained in a path of the elimination tree. Later, we will extend the algorithm to polynomials h that decompose into multiple paths of the elimination tree. Finally, we will prove the complexity bound from [Theorem 4.3](#).

Membership on a path. Consider the case where the elimination tree of the graph G is a path (i.e., it has only one leaf). Alternatively, we can assume that all the variables of h are contained in a path of the elimination tree. As before, let $h_C := h \bmod C$ denote the normal form with respect to chain C . Our radical ideal membership test is based on two simple ideas. First, we will check whether the polynomial $H(X) := \sum_C r_C h_C(X)$ is identically zero for some random coefficients $r_C \in \mathbb{K}$. Clearly, for sufficiently generic values of r_C , the polynomial $H(X)$ will be zero if and only if each h_C is zero. The second idea is that we evaluate $H(X)$ in some random points $\hat{x}_i \in \mathbb{K}$. Thus, we just need to check whether the scalar $H(\hat{X}) \in \mathbb{K}$ is zero. We illustrate how the algorithm works through the following example.

Example 4.6 (radical membership). Consider again the chordal network of [Figure 5](#). Let us verify that the polynomial $h(x)$ from [Figure 7](#) vanishes on its variety. We need to show that the reduction (normal form) of h by each chain of the network is zero. As in the case of counting solutions, we will achieve this by successively eliminating nodes. Note that the variables of h are $\{x_0, x_6, x_7, x_8, x_9\}$, which correspond to a path of the elimination tree. Thus, we restrict ourselves to the part of the network given by these variables, as shown in [Figure 7](#).

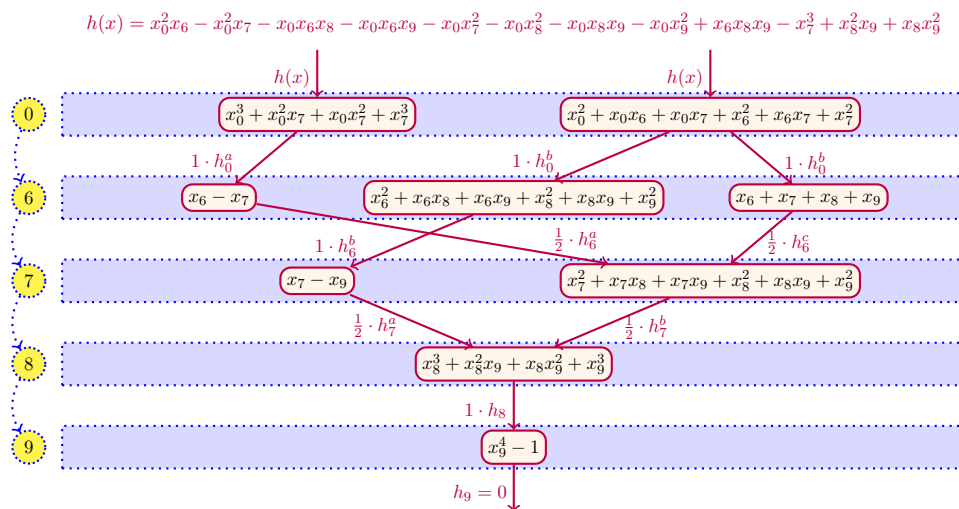


Figure 7. Sketch of the radical ideal membership test from [Example 4.6](#).

Let us start by processing the two nodes of rank 0. We have to compute the reduction of $h(x)$ modulo each of these nodes. Afterward, we will substitute x_0 in these reduced polynomials with a random value on \mathbb{K} ; in this case we choose $\hat{x}_0 = 1$. Let h_0^a, h_0^b be the polynomials obtained after the reduction and substitution, as shown in [Figure 7](#). These two polynomials will be sent to the adjacent rank 6 nodes.

Consider now a rank p node f_p that receives certain polynomials from its adjacent rank l nodes. We now perform a random linear combination of these incoming polynomials, then we reduce this linear combination modulo f_p , and finally we substitute x_p with a random value \hat{x}_p . For this example the linear combination will be an average, and the random points \hat{x}_p will be 1. Figure 7 indicates the polynomials received and output by each node. For instance, h_8 is obtained by reducing $\frac{1}{2}(h_7^a + h_7^b)$ modulo $f_8 = x_8^3 + x_8^2x_9 + x_8x_9^2 + x_9^3$ and then plugging in $\hat{x}_8 = 1$. The polynomials obtained with this procedure are shown below. Note that the last polynomial computed is zero, agreeing with the fact that $h(x)$ vanishes on the variety.

$$\begin{aligned} h_0^a &= x_6x_8x_9 - x_6x_8 - x_6x_9 + x_6 - x_7^3 - x_7^2 - x_7 + x_8^2x_9 - x_8^2 + x_8x_9^2 - x_8x_9 - x_9^2, \\ h_0^b &= -x_6^3 - x_6^2 + x_6x_8x_9 - x_6x_8 - x_6x_9 + x_8^2x_9 - x_8^2 + x_8x_9^2 - x_8x_9 - x_9^2, \\ h_6^a &= -x_7^3 - x_7^2 + x_7x_8x_9 - x_7x_8 - x_7x_9 + x_8^2x_9 - x_8^2 + x_8x_9^2 - x_8x_9 - x_9^2, \\ h_6^b &= -x_8^3 - x_8^2x_9 - x_8x_9^2 - x_9^3, \\ h_6^c &= x_7^3 + x_7^2(3x_8 + 3x_9 - 1) + x_7(3x_8^2 + 5x_8x_9 - x_8 + 3x_9^2 - x_9) + x_8^3 + 3x_8^2x_9 - x_8^2 + 3x_8x_9^2 - x_8x_9 + x_9^3 - x_9^2, \\ h_7^a &= h_7^b = -x_8^3 - x_8^2x_9 - x_8x_9^2 - x_9^3, \\ h_8 &= h_9 = 0. \end{aligned}$$

Algorithm 3 generalizes the procedure from the above example. Observing that each node f_l of the network has an associated polynomial $H(f_l)$, which is first reduced modulo f_l , then we substitute the value \hat{x}_l and finally we pass this polynomial to the adjacent nodes. Also note that we choose one random scalar \hat{x}_l for each variable, and one random scalar r_{lp} for each arc of the network.

Algorithm 3 Radical ideal membership

Input: Chordal network \mathcal{N} (triangular, square-free) and polynomial $h(x)$ such that all its variables are contained in a path of the elimination tree.

Output: True, if h vanishes on $\mathbf{V}(\mathcal{N})$. False, otherwise.

```

1: procedure RIDEALMEMBERSHIP( $\mathcal{N}, h$ )
2:    $x_m := \text{mvar}(h)$ 
3:   for  $f$  node of  $\mathcal{N}$  do
4:      $H(f) := (h \text{ if } \text{rk}(f) = m \text{ else } 0)$ 
5:   for  $l = 0 : n - 1$  do
6:      $\hat{x}_l := \text{random scalar}$ 
7:     for  $f_l$  node of  $\mathcal{N}$  of rank  $l$  do
8:        $H(f_l) := H(f_l) \bmod f_l$ 
9:       plug in  $\hat{x}_l$  in  $H(f_l)$ 
10:      for  $(f_l, f_p)$  arc of  $\mathcal{N}$  do
11:         $r_{lp} := \text{random scalar}$ 
12:         $H(f_p) := H(f_p) + r_{lp} H(f_l)$ 
13:   return (True if  $H(f_{n-1}) = 0$  for all rank  $n - 1$  nodes else False)

```

Correctness. We proceed to show the correctness of Algorithm 3. We will need a preliminary lemma and some new notation. For any l , let X^l denote the subtree of the elimination

tree consisting of x_l and all its descendants (e.g., X^{n-1} consists of all variables). For a rank l node f_l of the network, we will say that an f_l -subchain C_l is the subset of a chain C , with $f_l \in C$, restricted to nodes of rank i for some $x_i \in X^l$.

Lemma 4.7. *Let \mathcal{N} be a chordal network whose elimination tree is a path, and let $h \in \mathbb{K}[X]$. Let f_l be a rank l node of \mathcal{N} . In [Algorithm 3](#), the final value of $H(f_l)$ is given by plugging in the values $\hat{x}_1\hat{x}_2, \dots, \hat{x}_l$ in the polynomial*

$$\sum_{C_l} r_{C_l} h \bmod C_l,$$

where the sum is over all f_l -subchains C_l , and where r_{C_l} denotes the product of the random scalars r_{ij} along the subchain C_l .

Proof. See [Appendix A.2](#). ■

Theorem 4.8. *Let \mathcal{N} be a chordal network, triangular and square-free, and let q be a bound on the main degrees of its nodes. Let $h \in \mathbb{K}[X]$ be such that all its variables are contained in a path of the elimination tree. [Algorithm 3](#) behaves as follows:*

- *If h vanishes on $\mathbf{V}(\mathcal{N})$, it always returns “True.”*
- *If not, it returns “False” with probability at least $1/2$, assuming that the random scalars r_{lp}, x_l are chosen (independent, uniformly distributed) from some set $S \subseteq \mathbb{K}$ with $|S| \geq 2nq$.*

Proof. Denoting $h_C := h \bmod C$, [Lemma 4.7](#) tells us that [Algorithm 3](#) checks whether $\sum_C r_C h_C(\hat{X}) = 0$, where r_C is the product of all scalars r_{lp} along the chain C . If h vanishes on $\mathbf{V}(\mathcal{N})$, then each h_C is zero and thus the algorithm returns “True.” Assume now that h does not vanish on $\mathbf{V}(\mathcal{N})$, and thus at least one h_C is nonzero. Let R be the set of all random scalars r_{lp} used in the algorithm, which we now see as variables. Consider the polynomial

$$H(X, R) := \sum_C r_C(R) h_C(X),$$

and note that it is nonzero. Observe that the degree of $H(X, R)$ is at most nq , since $\deg(r_C) \leq n$ and $\deg(h_C) \leq n(q - 1)$. Using the Schwartz–Zippel lemma (see, e.g., [\[34, section 6.9\]](#)), the probability that H evaluates to zero for random values $r_{lp}, \hat{x}_l \in S$ is at most $nq/|S| \leq 1/2$. ■

Remark 4.9. The above theorem requires that \mathbb{K} contain sufficiently many elements. If necessary, we may consider a field extension $\mathbb{L} \supseteq \mathbb{K}$ and perform all computations over $\mathbb{L}[X]$.

Combining multiple paths. We now extend [Algorithm 3](#) to work for other polynomials h . Specifically, we assume that the polynomial can be written as $h = \sum_i h_i$, where the variables of each h_i belong to a path of the elimination tree. We let $x_{m_i} := \text{mvar}(h_i)$ denote the main variables, and we can assume that they are all distinct. We only need to make two simple modifications to [Algorithm 3](#).

- (i) Previously, we initialized the algorithm with nonzero values in a single rank (see [line 4](#)). We now initialize the algorithm in multiple ranks: $H(f_{m_i}) = h_i$ if $\text{rk}(f_{m_i}) = m_i$.
- (ii) When combining the incoming polynomials to a node f_p , we now take a random affine combination (i.e., $\sum_l r_{lp} H(f_l)$ for some scalars r_{lp} such that $\sum_l r_{lp} = 1$). Note that in

the example from Figure 7 we took the average of the incoming nodes, so this condition is satisfied.

The first modification is quite natural given the decomposition of the polynomial h . The second item is less intuitive, but it is simply a normalization to ensure that all polynomials h_i are scaled in the same manner. The correctness of this modified algorithm follows from the fact that Lemma 4.7 remains valid, as shown next.

Lemma 4.10. *Let \mathcal{N} be a chordal network, and let $h = \sum_i h_i \in \mathbb{K}[X]$ be such that the variables of each h_i are contained in a path of the elimination tree. With the above modifications to Algorithm 3, the final value of $H(f_l)$ is as stated in Lemma 4.7.*

Proof. See Appendix A.2. ■

Remark 4.11. Note that any h can be written as $h = \sum_{i=0}^{n-1} h_i$, where h_i corresponds to the terms with main variable x_i . Even when the elimination tree is a path, it is usually more efficient to decompose it in this manner and use Algorithm 3 with the above modifications.

Complexity. We finally proceed to prove the complexity bound from Theorem 4.3. We restrict ourselves to polynomials h that preserve the sparsity structure given by the chordal graph G . More precisely, we assume that the variables of each of the terms of h correspond to a clique of G , or, equivalently, that $h = \sum_l h_l$ for some $h_l \in \mathbb{K}[X_l]$. Naturally, we will use Algorithm 3 with the two modifications from above. The key idea to notice is that Algorithm 3 preserves chordality, as stated next.

Lemma 4.12. *Assume that in Algorithm 3 the initial values of $H(f_l)$ are such that $H(f_l) \subseteq \mathbb{K}[X_l]$ (where $\text{rk}(f_l) = l$). Then the same condition is satisfied throughout the algorithm.*

Proof. The update rule used in Algorithm 3 is of the form $H(f_p) := H(f_p) + r_{lp}\phi_l(\tilde{h}_l)$ for some $\tilde{h}_l \in \mathbb{K}[X_l]$, where ϕ_l denotes the functional that plugs in \hat{x}_l . Using Lemma 2.6, we have $\phi_l(\tilde{h}_l) \subseteq \mathbb{K}[X_l \setminus \{x_l\}] \subseteq \mathbb{K}[X_p]$. The result follows. ■

Proof of Theorem 4.3. We consider Algorithm 3 with the modifications (i) and (ii) from above. Note that the q -dominated condition allows us to bound the degrees of all polynomials computed in Algorithm 3. Furthermore, since chordality is preserved (Lemma 4.12), all polynomials will have at most q^κ terms. The complexity of the algorithm is determined by the cost of polynomial divisions and polynomial additions. Polynomial addition takes linear time in the number of terms, and it is performed once for each arc of the network. Thus, their total cost is $O(nW^2q^\kappa)$. As for polynomial division, $h \bmod f$ can be obtained in $O(|h| |f| \log |f|)$, where $|\cdot|$ denotes the number of terms [27]. Their total cost is $\tilde{O}(nWq^{2\kappa})$, since there is one operation per node of the network. ■

5. Monomial ideals. We already showed how to compute chordal network representations of some zero-dimensional ideals. Before proceeding to the general case, we will consider the special class of monomial ideals. Recall that an ideal is monomial if it is generated by monomials. Monomial ideals might have positive dimension, but their special structure makes their analysis particularly simple. As in Example 1.2, we will see that any monomial ideal admits a compact chordal network representation. We will also show how such a chordal network can be effectively used to compute its dimension, its equidimensional components, and its irreducible components. These methods will be later generalized to arbitrary polynomial

ideals.

5.1. Chordal triangularization. Algorithm 1 will be exactly the same for monomial ideals as in the zero-dimensional case. The only difference is that for the triangulation operations we need to specify the type of decomposition used, as explained now.

We will say that a set of monomials T is *triangular* if it consists of variables, i.e., $T = \{x_{i_1}, \dots, x_{i_m}\}$. It is well known that a monomial ideal is prime if and only if it is generated by variables. It is also known that the minimal primes of a monomial ideal are also monomial. It follows that any monomial ideal I decomposes as $\mathbf{V}(I) = \bigcup_T \mathbf{V}(T)$, where the union is over some triangular monomial sets T .

By using the above decomposition in each triangulation operation, chordal triangularization can now be applied to monomial ideals, as established in the proposition below. We point out that even though this decomposition seems quite different from that of section 3.1, both are special instances of a more general theory that will be discussed in section 6.1.

Proposition 5.1. *Let F be a set of monomials supported on a chordal graph G . Algorithm 1 computes a G -chordal network, whose chains give a triangular decomposition of F .*

Proof. Proving that the variety is preserved in the algorithm is essentially the same as for the chordally zero-dimensional case (Lemma 3.13). It is straightforward to see that the chains of the output are triangular (i.e., they consist of variables). ■

Example 5.2. Consider the ideal $I = \langle x_0x_1, x_0x_2, x_0x_3, x_1x_2, x_1x_4, x_2x_5, x_3x_4, x_3x_5, x_4x_5 \rangle$. The result of chordal triangularization is shown on the left of Figure 8.




Figure 8. Chordal network from Example 5.2 and its top-dimensional part.

As in the chordally zero-dimensional case, we can also prove that the complexity is linear in n when the treewidth is bounded.

Theorem 5.3. *Let F be a set of monomials supported on a chordal graph G of clique number κ . Then F can be represented by a triangular chordal network with at most $n2^\kappa$ nodes, which can be computed in time $O(n2^{O(\kappa)})$.*

Proof. Note that after the l th triangulation round we will have at most 2^κ rank l nodes, since the triangular monomial sets in $\mathbb{K}[X_l]$ are in bijection with the subsets of X_l . A similar argument proves that the width of the network is bounded by 2^κ after an elimination round, and thus throughout the algorithm. The cost of computing a triangular decomposition in $\mathbb{K}[X_l]$ is polynomial in $2^{|X_l|}$, since we can simply enumerate over all possible triangular monomial

sets. Thus, the cost of all triangulation operations is $O(nW 2^{O(\kappa)}) = O(n 2^{O(\kappa)})$. The cost of the elimination and merging operations is negligible. ■

5.2. Computing with chordal networks. Let \mathcal{N} be a chordal network representation of a monomial ideal I . We will show how to effectively use \mathcal{N} to solve the following problems:

Dimension. Determine the dimension of I .

Top-dimensional part. Describe the top-dimensional part of $\mathbf{V}(I)$.

Irreducible components. Determine the minimal primes of I .

The above problems can be shown to be hard in general by using the correspondence between minimal vertex covers of a graph and the irreducible components of its edge ideal (see [Example 1.2](#)). We will see that, given the chordal network, the first two problems can be solved in linear time with a dynamic program. The third one is much more complicated, since we need to enumerate over all chains of the network to verify if they are minimal. In order to do this efficiently, we will need to address the following problems.

Dimension count. Classify the number of chains C of \mathcal{N} according to its dimension.

Isolate dimension d . Enumerate all chains C of \mathcal{N} such that $\dim(\mathbf{V}(C)) = d$.

We proceed to solve each of the problems from above. To simplify the exposition, we will assume for this section that the elimination tree is a path, but it is not difficult to see that all these methods will work for arbitrary chordal networks.

Dimension. Let us see that it is quite easy to compute the dimension of $\mathbf{V}(\mathcal{N})$. Since the variety $\mathbf{V}(T)$ of a triangular monomial set is a linear space, its dimension is $\dim(\mathbf{V}(T)) = n - |T|$. Therefore, $\dim(\mathbf{V}(\mathcal{N})) = n - \min_C |C|$, where the minimum is taken over all chains of the network. Note that we ignore the zero entries of C . In particular, for the network in [Figure 8](#) we have $\dim(\mathbf{V}(\mathcal{N})) = 6 - 4 = 2$.

We reduced the problem to computing the smallest cardinality of a chain of \mathcal{N} . This can be done using a simple dynamic program, which is quite similar to that in [Algorithm 2](#). For each node f_l we save the value $\ell(f_l)$ corresponding to the length of the shortest chain up to level l . For an arc (f_l, f_p) with $f_p \neq 0$, the update rule is simply $\ell(f_p) := \min(\ell(f_p), 1 + \ell(f_l))$. It follows that we can compute in linear time the dimension of $\mathbf{V}(\mathcal{N})$.

Top-dimensional part. We can get a chordal network \mathcal{N}_{top} describing its top-dimensional part by modifying the procedure that computes the dimension. Indeed, assume that for some arc (f_l, f_p) we have $\ell(f_p) < 1 + \ell(f_l)$, and thus the update $\ell(f_p) := \min(\ell(f_p), 1 + \ell(f_l))$ is not needed. This means that the arc (f_l, f_p) is unnecessary for the top-dimensional component. By pruning the arcs of \mathcal{N} in such a manner, we obtain the wanted network \mathcal{N}_{top} .

Example 5.4. Let \mathcal{N} be the network on the left of [Figure 8](#). Note that \mathcal{N} has nine chains; two of them are $C_1 = (x_1, x_2, x_3, x_5)$, $C_2 = (x_0, x_1, x_2, x_3, x_5)$, of dimensions 2 and 1. By pruning some arcs, we obtain its highest dimensional part \mathcal{N}_{top} , shown on the right of [Figure 8](#). This network \mathcal{N}_{top} only has six chains; note that C_2 is not one of them. In this case neither of the chains removed was minimal (e.g., $C_2 \supsetneq C_1$), so that $\mathbf{V}(\mathcal{N}) = \mathbf{V}(\mathcal{N}_{top})$. Thus, both \mathcal{N} and \mathcal{N}_{top} are valid chordal network representations of the ideal from [Example 5.2](#), although the latter is preferred since all its chains are minimal. Similarly, the network from [Figure 2](#) was obtained by using chordal triangularization and then computing its highest dimensional part.

Irreducible components. Chordal triangularization can also aid in computing the minimal primes of an ideal (geometrically, the irreducible components). In the monomial case, any chain of \mathcal{N} defines a prime ideal, and thus we only need to determine which chains are minimal with respect to containment. In some cases it is enough to prune certain arcs of the network (e.g., Figure 8), but this is not always possible.

Unfortunately, we do not know a better procedure than simply iterating over all chains of the network, checking for minimality. Nonetheless, we can make this method much more effective by proceeding in order of decreasing dimension. This simple procedure is particularly efficient when we are only interested in the minimal primes of high dimension, as will be seen in section 7.1. In the remainder of the section we will explain how to enumerate the chains by decreasing dimension (this is precisely the dimension isolation problem).

Dimension count. Classifying the number of chains according to its dimension can be done with a dynamic program very similar to that used for computing the dimension. As discussed above, the dimension of a chain is simply given by its cardinality. For a rank l node f_l of the network and for any $0 \leq k \leq l + 1$, let $c_k(f_l)$ denote the number of chains of the network (up to level l) with cardinality exactly k . Then for an arc (f_l, f_p) with $f_p \neq 0$ the update rule is simply $c_k(f_p) := c_k(f_p) + c_{k-1}(f_l)$.

Dimension isolation. For simplicity of exposition we only describe how to produce one chain C of dimension d , but it is straightforward to then generate all of them. As in Example 4.2, we follow a bottom-up strategy, successively adding nodes to the chain. We first need to choose a rank $n - 1$ node f_{n-1} that belongs to at least one chain of dimension d . Using the values $c_k(f_l)$ from above, we can choose any f_{n-1} for which $c_{n-d}(f_{n-1}) \geq 1$. Assuming that we chose some $f_{n-1} \neq 0$, we now need to find an adjacent rank $n - 2$ node f_{n-2} such that $c_{n-d-1}(f_{n-2}) \geq 1$. It is clear how to continue.

6. The general case. We finally proceed to compute chordal network representations of arbitrary polynomial ideals. We will also see how the different chordal network algorithms developed earlier (e.g., radical ideal membership, isolating the top-dimensional component) have a natural extension to this general setting.

6.1. Regular chains. The theory of triangular sets for positive-dimensional varieties is more involved; we refer the reader to [17, 36] for an introduction. We now present the concept of regular chains, which is at the center of this theory.

A set of polynomials $T \subseteq \mathbb{K}[X] \setminus \mathbb{K}$ is a *triangular set* if its elements have distinct main variables. Let h be the product of the initials (Definition 3.3) of the polynomials in T . The geometric object associated to T is the *quasi-component*

$$\mathbf{W}(T) := \mathbf{V}(T) \setminus \mathbf{V}(h) \subseteq \overline{\mathbb{K}^n}.$$

The attached algebraic object is the *saturated ideal*

$$\text{sat}(T) := \langle T \rangle : h^\infty = \{f \in \mathbb{K}[X] : h^N f \in \langle T \rangle \text{ for some } N \in \mathbb{N}\}.$$

Note that $\mathbf{V}(\text{sat}(T)) = \overline{\mathbf{W}(T)}$, where the closure is in the Zariski topology.

Polynomial pseudodivision is a basic operation in triangular sets. Let f, g be polynomials of degrees d, e in $x := \text{mvar}(g)$. The basic idea is to see f, g as univariate polynomials in x (with coefficients in $\mathbb{K}[X \setminus \{x\}]$), and in order that we can always divide f by g , we first multiply by some power of $\text{init}(g)$. Formally, the *pseudoremainder* of f by g is $\text{prem}(f, g) := f$ if $d < e$, and otherwise $\text{prem}(f, g) := \text{init}(g)^{d-e+1} f \bmod (g)$. Pseudodivision can be extended to triangular sets in the natural way. The *pseudoremainder* of f by $T = \{t_1, \dots, t_k\}$, where $\text{mvar}(t_1) > \dots > \text{mvar}(t_k)$, is

$$\text{prem}(f, T) = \text{prem}(\dots(\text{prem}(f, t_1) \dots, t_k).$$

Definition 6.1. A regular chain is a triangular set T such that for any polynomial f

$$f \in \text{sat}(T) \iff \text{prem}(f, T) = 0.$$

Remark 6.2. Note that a zero-dimensional triangular set (**Definition 3.3**) is a regular chain, since pseudoreduction coincides with Gröbner bases reduction.

Regular chains have very nice algorithmic properties. In particular, they are always consistent (i.e., $\mathbf{W}(T) \neq 0$), and furthermore $\dim(\mathbf{W}(T)) = n - |T|$. **Table 1** summarizes some of these properties, comparing them with Gröbner bases.

Table 1
Gröbner bases versus regular chains.

	Gröbner basis (\mathcal{G})	Regular chain (T)
Geometric object	$\mathbf{V}(\mathcal{G})$	$\mathbf{W}(T)$
Algebraic object	$\langle \mathcal{G} \rangle$	$\text{sat}(T)$
Feasible	if $1 \notin \mathcal{G}$	always
Ideal membership	remainder = 0	pseudoremainder = 0
Dimension	from Hilbert series	$n - T $
Elimination ideal	$\mathcal{G}_{\text{lex}} \cap \mathbb{K}[x_1, \dots, x_{n-1}]$	$T \cap \mathbb{K}[x_1, \dots, x_{n-1}]$

Definition 6.3. A triangular decomposition of a polynomial set F is a collection \mathcal{T} of regular chains, such that $\mathbf{V}(F) = \bigcup_{T \in \mathcal{T}} \mathbf{W}(T)$.

Remark 6.4. There is a weaker notion of decomposition that is commonly used: \mathcal{T} is a *Kalkbrener* triangular decomposition if $\mathbf{V}(F) = \bigcup_{T \in \mathcal{T}} \overline{\mathbf{W}(T)}$.

Example 6.5. Let $F = \{x_0x_3 - x_1x_2, x_2x_5 - x_3x_4, x_4x_7 - x_5x_6\}$ consist of the adjacent minors of a 2×4 matrix. It can be decomposed into eight regular chains:

$$\begin{aligned} & (x_0x_3 - x_1x_2, x_2x_5 - x_3x_4, x_4x_7 - x_5x_6), (x_0x_3 - x_1x_2, x_4, x_5), (x_2, x_3, x_4x_7 - x_5x_6), \\ & (x_1, x_3, x_4, x_5), (x_1, x_3, x_5, x_7), (x_2, x_3, x_5, x_7), (x_2, x_3, x_6, x_7), (x_0x_3 - x_1x_2, x_2x_5 - x_3x_4, x_6, x_7). \end{aligned}$$

Note that the first three triangular sets (first line) have dimension $8 - 3 = 5$. Observe that the quasi-components $\mathbf{W}(T)$ of these three sets do not cover the points for which $x_3 = x_7 = 0$, which is why we need the remaining five sets. However, these three triangular sets alone give a Kalkbrener decomposition of the variety.

6.2. Regular systems. In the study of triangular sets, it is useful to consider systems of polynomials containing both equations $\{f_i(x) = 0\}_i$ and *inequations* $\{h_j(x) \neq 0\}_j$. Following the notation of [36], we say that a *polynomial system* $\mathfrak{F} = (F, H)$ is a pair of polynomial sets $F, H \subseteq \mathbb{K}[X]$, and its associated geometric object is the quasi-variety

$$\mathbf{Z}(\mathfrak{F}) := \{x \in \overline{\mathbb{K}}^n : f(x) = 0 \text{ for } f \in F, h(x) \neq 0 \text{ for } h \in H\}.$$

For instance, the quasi-component $\mathbf{W}(T)$ of a triangular set is the quasi-variety of the polynomial system $(T, \text{init}(T))$, where $\text{init}(T)$ is the set of initials of T .

For a polynomial system $\mathfrak{F} = (F, H)$ we denote by $\text{elim}_p(\mathfrak{F})$ the polynomial system $(\text{elim}_p(F), \text{elim}_p(H))$. We also denote by $\mathfrak{F}_1 + \mathfrak{F}_2$ the concatenation of two polynomial systems, i.e., $(F_1, H_1) + (F_2, H_2) := (F_1 \cup F_2, H_1 \cup H_2)$.

Definition 6.6. A regular system is a pair $\mathfrak{T} = (T, U)$ such that T is triangular and for any $0 \leq k < n$ the following hold:

- (i) Either $T^{(k)} = \emptyset$ or $U^{(k)} = \emptyset$, where the superscript $^{(k)}$ denotes the polynomials with main variable x_k .
- (ii) $\text{init}(f)(\hat{x}^{k+1}) \neq 0$ for any $f \in T^{(k)} \cup U^{(k)}$ and $\hat{x}^{k+1} \in \mathbf{Z}(\text{elim}_{k+1}(\mathfrak{T})) \subseteq \overline{\mathbb{K}}^{n-k-1}$.

A regular system is square-free if the polynomials $f(x_k, \hat{x}^{k+1}) \subseteq \mathbb{K}[x_k]$ are square-free for any $f \in T^{(k)} \cup U^{(k)}$ and any $\hat{x}^{k+1} \in \mathbf{Z}(\text{elim}_{k+1}(\mathfrak{T})) \subseteq \overline{\mathbb{K}}^{n-k-1}$.

For a regular system (T, U) the set T is a regular chain, and, conversely, for a regular chain T there is some U such that (T, U) is a regular system [35]. Wang showed how to decompose any polynomial system in characteristic zero into (square-free) regular systems [35, 36].

Definition 6.7. A triangular decomposition of a polynomial system \mathfrak{F} is a collection \mathcal{T} of regular systems, such that $\mathbf{Z}(\mathfrak{F}) = \bigcup_{\mathfrak{T} \in \mathcal{T}} \mathbf{Z}(\mathfrak{T})$.

Remark 6.8 (binomial ideals). Consider a polynomial system $\mathfrak{F} = (F, U)$ such that F consists of binomials (two terms) and $H = \{x_{i_1}, \dots, x_{i_m}\}$ consists of variables. We can decompose \mathfrak{F} into regular systems $\mathfrak{T} = (T, U)$ that preserve the binomial structure. Assume first that $H = \{x_0, \dots, x_{n-1}\}$ contains all variables. Equivalently, we are looking for the zero set of F on the torus $(\overline{\mathbb{K}} \setminus \{0\})^n$. It is well known that F can be converted to (binomial) triangular form T by computing the Hermite normal form of the matrix of exponents [31, section 3.2], and the inequations U correspond to the nonpivot variables. For an arbitrary H , we can enumerate over the choices of nonzero variables.

6.3. Chordal triangularization. Algorithm 1 extends to the positive-dimensional case in the natural way, although with one important difference: the nodes of the chordal network will be polynomial systems, i.e., pairs of polynomial sets. We now describe the modifications of the main steps of the algorithm:

Initialization. The nodes of the induced G -chordal network are now of the form $\mathfrak{F}_l = (F_l, H_l)$, where $F_l = F \cap \mathbb{K}[X_l]$ and $H_l = \emptyset$.

Triangulation. For a node \mathfrak{F}_l we decompose it into regular systems \mathfrak{T} , and we replace \mathfrak{F}_l with a node for each of them.

Elimination. Let \mathfrak{F}_l be the rank l node we will eliminate, and let \mathfrak{F}_p be an adjacent rank p node. Then we create a rank p node $\mathfrak{F}'_p := \mathfrak{F}_p + \text{elim}_p(\mathfrak{F}_l)$.

Termination. After all triangulation/elimination operations, we may remove the inequations from the nodes of the network, i.e., replace $\mathfrak{F} = (F, H)$ with F .

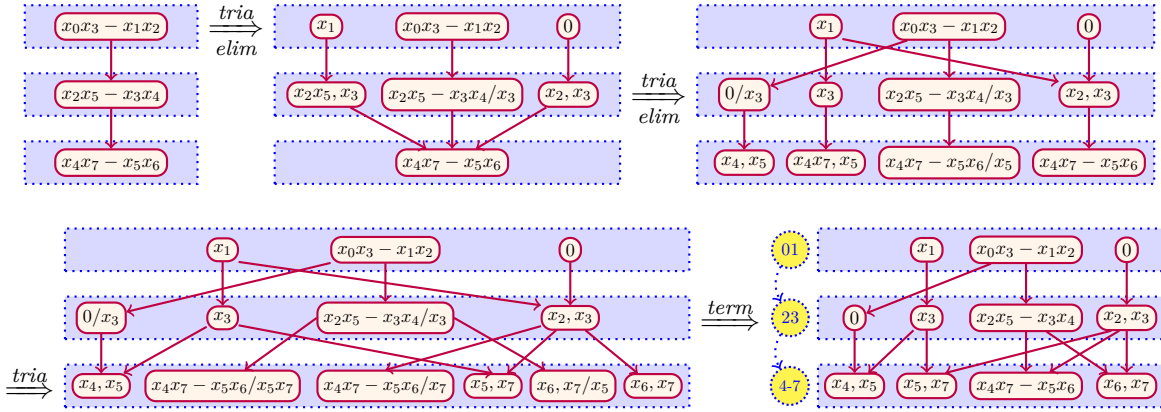


Figure 9. Chordal triangularization from Example 6.9.

Example 6.9. Figure 9 illustrates the chordal triangularization algorithm for the polynomial set F from Example 6.5. The nodes of the chordal network are polynomial systems $\mathfrak{F} = (F, H)$, which we represent in the figure as (F/H) . Note that in the termination step, after all triangulation/elimination operations, we remove the inequations to simplify the network. The final network has eight chains, which coincide with the triangular decomposition from Example 6.5.

We can now compute chordal network representations of arbitrary systems.

Theorem 6.10. Let $F \subseteq \mathbb{K}[X]$ be supported on a chordal graph G . With the above modifications, Algorithm 1 computes a G -chordal network \mathcal{N} , whose chains give a triangular decomposition of F . Furthermore, this decomposition is square-free if all triangulation operations are square-free.

Proof. See Appendix A.3. ■

Remark 6.11. We have noticed that chordal triangularization is quite efficient for binomial ideals. Remark 6.8 partly explains this observation. However, we do not yet know whether it will always run in polynomial time when the treewidth is bounded.

6.4. Computing with chordal networks. We just showed how to compute chordal network representations of arbitrary polynomial systems. We now explain how to extend the chordal network algorithms from section 4 and section 5.2 to the general case.

Elimination. Since regular chains possess the same elimination property as lexicographic Gröbner bases, the approach from section 4.1 works in the same way.

Radical ideal membership. Algorithm 3 extends to the positive-dimensional case simply by replacing polynomial division with pseudodivision. Note that we require a square-free chordal network, which can be computed as explained in Theorem 6.10.

Dimension and equidimensional components. The dimension of a regular chain T is $n - |T|$, which is the same as for the monomial case. Thus, we can compute the dimension as in section 5.2. Similarly, we can compute a chordal network describing the highest dimensional component, and also isolate any given dimension of the network.




Figure 10. Top-dimensional part of the chordal network from Figure 9.

Example 6.12. Figure 10 shows the highest dimensional part of the chordal network from Figure 9. This network only has three chains, which give a Kalkbrener decomposition of the variety (see Example 6.5). Likewise, the chordal network from Figure 3 gives a Kalkbrener triangular decomposition of the ideal of adjacent minors of a 2×7 matrix.

Irreducible components. Unlike the monomial case, the chains of an arbitrary chordal network \mathcal{N} will not necessarily define prime ideals (see section 3.4). However, in some interesting cases this will be the case, thanks to the following well-known property.

Theorem 6.13. Let $T = \{t_1, \dots, t_k\}$ be a regular chain. Assume that $\text{mdeg}(t_i) = 1$ for $1 \leq i < k$ and that t_k is an irreducible polynomial. Then $\text{sat}(T)$ is a prime ideal.

Proof. This follows from [36, Thm. 6.2.14]. ■

In particular, note that all chains of the chordal network from Figure 3 are of this form. We will see in section 7.3 that the same holds for other families of ideals. Assume now that all chains of the network define a prime ideal. A plausible strategy to compute all minimal primes (or only the high dimensional ones) is as follows:

- (i) Iterate over all chains T of the network in order of decreasing dimension.
- (ii) For a chain C , and a minimal prime I' previously found, determine whether $I' \subseteq \text{sat}(C)$ by checking whether $\text{prem}(f, C) = 0$ for each generator f of I' .
- (iii) If $I := \text{sat}(C)$ does not contain any previously found prime, compute generators for I by using Gröbner bases. We have a new minimal prime.

7. Examples. We conclude this paper by exhibiting some examples of our methods. We implemented our algorithms in Sage [30] using Maple’s library Epsilon [37] for triangular decompositions, and Singular [12] for Gröbner bases. The experiments are performed on an i7 PC with 3.40GHz, 15.6 GB RAM, running Ubuntu 14.04.

7.1. Commuting birth and death ideal. We consider the binomial ideal I^{n_1, \dots, n_k} from [14]. This ideal models a generalization of the one-dimensional birth-death Markov process to higher dimensional grids. In [14] it is given a parametrization of its top-dimensional component, as well as the primary decomposition of some small cases. In [18] Kahle uses his Macaulay2 package Binomials, specialized in binomial ideals, to compute primary decompositions of

larger examples. We now show how our methods can greatly surpass Kahle's methods in computing the irreducible decomposition when the treewidth is small.

We focus on the case of a two-dimensional grid:

$$I^{n_1, n_2} = \langle U_{i,j}R_{i,j+1} - R_{i,j}U_{i+1,j}, D_{i,j+1}R_{i,j} - R_{i,j+1}D_{i+1,j+1}, \\ D_{i+1,j+1}L_{i+1,j} - L_{i+1,j+1}D_{i,j+1}, U_{i+1,j}L_{i+1,j+1} - L_{i+1,j}U_{i,j} \rangle_{0 \leq i < n_1, 0 \leq j < n_2}.$$

We let the parameter n_1 take values between 1 to 100, while n_2 is either 1 or 2. [Table 2](#) shows the time used by [Algorithm 1](#) for different values of n_1, n_2 . Observe that, for small values of n_2 , our methods can handle very high values of n_1 thanks to our use of chordality. For comparison, we note that even for the case $n_1 = 10, n_2 = 1$ [Singular's](#) Gröbner basis algorithm (grevlex order) did not terminate within 20 hours of computation. Similarly, neither [Epsilon](#) [37] nor [RegularChains](#) [25] was able to compute a triangular decomposition of $I^{10,1}$ within 20 hours.

Table 2

Time required by chordal triangularization on ideals I^{n_1, n_2} . No other software we tried [12, 18, 25, 37] can solve these problems.

n_1	20	40	60	80	100
$n_2 = 1$	0:00:45	0:02:16	0:04:03	0:06:28	0:09:13
$n_2 = 2$	0:36:07	1:59:24	3:30:33	6:15:25	9:00:52

We now consider the computation of the irreducible components of the ideal $I^{n_1, 1}$. We follow the strategy described after [Theorem 6.13](#), using [Sage's](#) default algorithm to compute saturations. [Table 3](#) compares this strategy (including the time of [Algorithm 1](#)) with the algorithm from [Binomials](#) [18]. It can be seen that our method is more efficient. In particular, for the ideal $I^{7,1}$ Kahle's algorithm did not finish within 60 hours of computation.

Table 3

Irreducible components of the ideals $I^{n_1, 1}$.

n_1	1	2	3	4	5	6	7
#components	3	11	40	139	466	1528	4953
time							
ChordalNet	0:00:00	0:00:01	0:00:04	0:00:13	0:02:01	0:37:35	12:22:19
Binomials	0:00:00	0:00:00	0:00:01	0:00:12	0:03:00	4:15:36	-

Comparing [Table 2](#) with [Table 3](#) it is apparent that computing a triangular chordal network representation is considerably simpler than computing the irreducible components. Nonetheless, if we are only interested in the high dimensional components, the complexity can be significantly improved. Indeed, in [Table 4](#) we can see how we can very efficiently compute all components of the seven highest dimensions.

7.2. Lattice walks. We now show a simple application of our radical membership test. We consider the lattice reachability problem from [13] (see also [22]). Given a set of vectors $\mathcal{B} \subseteq \mathbb{Z}^n$, construct a graph with vertex set \mathbb{N}^n in which $u, v \in \mathbb{N}^n$ are adjacent if $u - v \in \pm \mathcal{B}$. The problem is to decide whether two vectors $s, t \in \mathbb{N}^n$ are in the same connected component

Table 4
High dimensional irreducible components of the ideals $I^{n_1,1}$.

n_1	Highest 5 dimensions					Highest 7 dimensions			
	20	40	60	80	100	10	20	30	40
#comps	404	684	964	1244	1524	2442	5372	8702	12432
time	0:01:07	0:04:54	0:15:12	0:41:52	1:34:05	0:05:02	0:41:41	3:03:29	9:53:09

of the graph. This problem is equivalent to an ideal membership problem for a certain binomial ideal $I_{\mathcal{B}}$ [13]. Therefore, our radical membership test can be used to prove that s, t are not in the same connected component (but it may fail to prove the converse). We consider the following sample problem.




Figure 11. Illustration of the card problem using 5 decks.

Problem 7.1. There are n card decks organized on a circle, as shown in Figure 11. Given any four consecutive decks we are allowed to move the cards as follows: we may take one card from each of the inner decks and place them in the outer decks (one in each), or we may take one card from the outer decks and place them in the inner decks. Initially the number of cards in the decks is $1, 2, \dots, n$. Is it possible to reach a state where the number of cards in the decks is reversed³ (i.e., the i th deck has $n - i + 1$ cards)?

The above problem is equivalent to determining whether $f_n \in I_n$, where

$$f_n := x_0 x_1^2 x_2^3 \cdots x_{n-1}^n - x_0^n x_1^{n-1} \cdots x_{n-1}, \quad I_n := \{x_i x_{i+3} - x_{i+1} x_{i+2} : 0 \leq i < n\},$$

and where the indices are taken modulo n . Table 5 compares our method against Singular’s Gröbner basis (grevlex order) and Epsilon’s triangular decomposition. Even though the ideal I_n is not radical, in all experiments performed we obtained the right answer. Note that the complexity of our method is almost linear. This contrasts with the exponential growth of both Singular and Epsilon, which did not terminate within 20 hours for the cases $n = 30$ and $n = 45$. We do not include timings for Binomials and RegularChains since they are both slower than Singular and Epsilon.

7.3. Finite state diagram representation. One of the first motivations in this paper was the very nice chordal network representation of the irreducible components of the ideal of

³A combinatorial argument proves that this is only possible if all prime divisors of n are at least 5. However, this argument does not generalize to other choices of the final state (e.g., we cannot reach a state where the number of cards is $2, 1, 3, 4, 5, \dots, n$ for any n).

Table 5
Time (seconds) to test (radical) ideal membership on the ideals I_n .

n	5	10	15	20	25	30	35	40	45	50	55
ChordalNet	0.7	3.0	8.5	14.3	21.8	29.8	37.7	48.2	62.3	70.6	84.8
Singular	0.0	0.0	0.2	17.9	1036.2	-	-	-	-	-	-
Epsilon	0.1	0.2	0.4	2.0	54.4	160.1	5141.9	17510.1	-	-	-
Test result	true	false	false	false	true	false	true	false	false	false	true

adjacent minors of a $2 \times n$ matrix. We will see now that similar chordal network representations exist for other determinantal ideals.

First, notice that the chordal network in Figure 3 has a simple pattern. Indeed, there are three types of nodes $A_i = \{x_{2i}x_{2i+3} - x_{2i+1}x_{2i+2}\}$, $B_i = \{0\}$, $C_i = \{x_{2i}, x_{2i+1}\}$, and we have some valid transitions: $A_i \rightarrow \{A_{i+1}, B_{i+1}\}$, $B_i \rightarrow \{C_{i+1}\}$, $C_i \rightarrow \{A_{i+1}, B_{i+1}\}$. This transition pattern is represented in the state diagram shown in Figure 12a. Following the convention from automata theory, we mark the initial states with an incoming arrow and the terminal states with a double line.

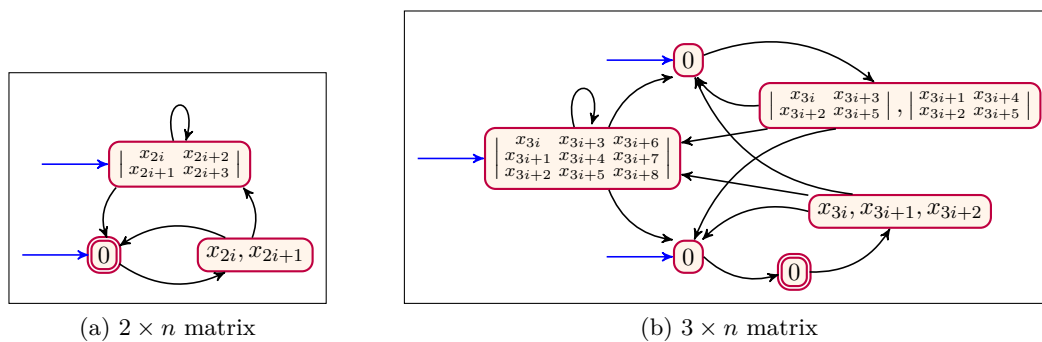


Figure 12. State diagrams for ideals of adjacent minors of a matrix.

We can also consider the ideal of 3×3 adjacent minors of a $3 \times n$ matrix. As seen in Figure 12b, a very similar pattern arises. In order to make sense of such a diagram, let us think of how to generate a $3 \times n$ matrix satisfying all these minor constraints. Let $v_1, \dots, v_n \in \mathbb{K}^3$ denote the column vectors. Given v_{i+1}, v_{i+2} , we can generate v_i as follows: it can be the zero vector, or it can be a multiple of v_{i+1} , or it can be a linear combination of v_{i+1}, v_{i+2} . These three choices correspond to the three main states shown in the diagram. Note now that if v_i is the zero vector, then we can ignore it when we generate v_{i-1} and v_{i-2} . This is why in order to reach the state $(x_{3i}, x_{3i+1}, x_{3i+2})$ we have to pass two trivial states. Similarly, if v_i is parallel to v_{i+1} , then we can ignore v_{i+1} when we generate v_{i-1} .

It is easy to see that the above reasoning generalizes if we consider the adjacent minors of a $k \times n$ matrix. Therefore, for any fixed k , the ideal of adjacent minors of a $k \times n$ matrix has a finite state diagram representation (and thus it has a chordal network representation of linear size). Since the nodes of the network are given by minors, then all chains of the network are of the form of Theorem 6.13. Thus, the decomposition obtained is into irreducible components.

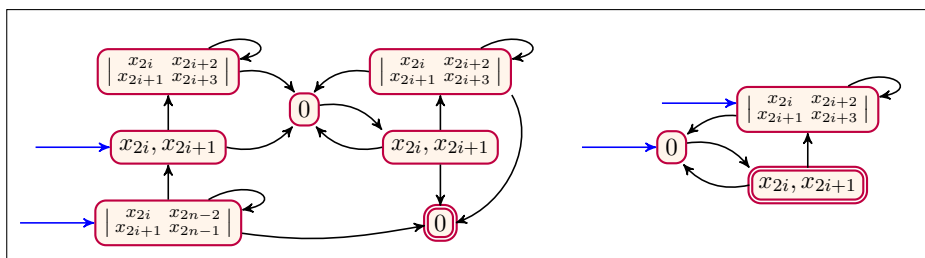


Figure 13. State diagram for the ideal of cyclically adjacent minors of a $2 \times n$ matrix.

Many other families of ideals admit a simple state diagram representation—for instance, the ideal generated by the n cyclically adjacent minors of a $2 \times n$ matrix (see Figure 13). Interestingly, this chordal network has two equidimensional components. Similarly, the ideal of (cyclically) adjacent permanent minors has a finite state diagram representation. We can also easily provide families of zero-dimensional problems with such a property (e.g., Figure 1), since they often admit a chordal network of linear size (Remark 3.23). Similar reasoning applies for monomial ideals. It is natural to ask for further examples of this behavior.

Question 7.2. Characterize interesting families of ideals (parametrized by n) whose triangular decomposition admits a finite state diagram representation (and thus has a chordal network representation of size $O(n)$).

Remark 7.3. The class of binomial edge ideals [16] is a natural starting point for this question, given that it generalizes both the ideal of adjacent minors (Figure 12a) and cyclically adjacent minors (Figure 13) of a $2 \times n$ matrix.

Appendix A. Additional proofs.

A.1. Proofs from section 3.

Proof of Lemma 3.10. Let $m < n$ be such that X_m is a maximal clique, and consider a rank m node $F_m \subseteq \mathbb{K}[X_m]$ to which we will apply a triangulation operation. Also let $F'_m := F_m \cap \mathbb{K}[X_m]$ be the unique initial node of rank m . By assumption, F'_m is zero-dimensional. Note that when we create a new node of rank m in an elimination operation, we copy the equations from a previous rank m node. In particular, we must have that $F'_m \subseteq F_m$, and therefore F_m is also zero-dimensional. This proves the lemma for this case.

Consider now some $p < n$ such that X_p is not maximal, which means that x_p is not a leaf of the elimination tree. Since X_p is not maximal, there is a child x_l of x_p such that $X_l = X_p \cup \{x_l\}$. By induction, we may assume that the lemma holds for all nodes of rank l . Consider a rank p node $F_p \subseteq \mathbb{K}[X_p]$ that we want to triangulate, and let F_l of rank l be adjacent to F_p . Let F'_l be the same rank l node, but before the l th elimination round. By induction, $F'_l \subseteq \mathbb{K}[X_l]$ is zero-dimensional. Therefore, $\text{elim}_p(F'_l) \subseteq \mathbb{K}[X_l \setminus \{x_l\}] = \mathbb{K}[X_p]$ is also zero-dimensional, and as $\text{elim}_p(F'_l) \subseteq F_p$, we conclude that F_p is zero-dimensional. ■

Lemma A.1. Let $X_1, X_2 \subseteq X$, and let $I_1 \subseteq \mathbb{K}[X_1]$, $I_2 \subseteq \mathbb{K}[X_2]$ be radical zero-dimensional ideals. Then $I_1 + I_2 \subseteq \mathbb{K}[X_1 \cup X_2]$ is also radical and zero-dimensional.

Proof. This is a direct consequence of the following known fact (see, e.g., [31, Thm. 2.2]):

an ideal $I \subseteq \mathbb{K}[X]$ is radical and zero-dimensional if and only if for any $x_i \in X$ there is a nonzero square-free polynomial $f \in I \cap \mathbb{K}[x_i]$. ■

Proof of Proposition 3.14. For any l , let X^l denote the subtree of the elimination tree consisting of x_l and all its descendants. For a chordal network \mathcal{N} , we will say that an l -subchain C_l is the subset of a chain C restricted to nodes with rank i for some $x_i \in X^l$. Note that any chain is also an $(n-1)$ -chain. Thus, it suffices to show that every l -subchain is radical after the l th triangulation round in Algorithm 1, and we proceed to show it by induction.

If x_l is a leaf in the elimination tree, then any l -subchain is just the output of a triangulation operation, and thus it is radical. Assume that the result holds for all $l < p$. Let T_p be a rank p node obtained after the p th triangulation round. Let C be a p -subchain containing T_p ; we want to show that $\langle C \rangle$ is radical. Let x_{l_1}, \dots, x_{l_k} be the children of x_p . For each l_j , let C_{l_j} be the l_j -subchain obtained by restricting C to ranks in X^{l_j} . Also let C'_{l_j} be the same l_j -subchain, but before the l_j th elimination round. Observe that

$$\langle C \rangle = \langle T_p \rangle + \sum_j \langle C'_{l_j} \rangle.$$

Note that $\langle T_p \rangle$ is zero-dimensional and radical, and by induction the same holds for each $\langle C'_{l_j} \rangle$. It follows from Lemma A.1 that $\langle C \rangle$ is radical. ■

Proof of Proposition 3.20. It was shown in [15] that the complexity of Buchberger's algorithm is $q^{O(k)}$ if the equations $x_i^q - x_i$ are present, and the same analysis works for any q -dominated ideal. Given a Gröbner basis, the LexTriangular algorithm [24] computes a triangular decomposition in time $D^{O(1)}$, where $D \leq q^k$ is the number of standard monomials. For irreducible (or square-free) decompositions, we can reduce the problem to the univariate case by using a rational univariate representation [29] (here we need that \mathbb{K} contains sufficiently many elements). This representation can also be obtained in $D^{O(1)}$. Since the complexity of univariate (square-free) factorization [20] is polynomial in the degree (D), the result follows. ■

Proof of Lemma 3.22. Let us see that the result holds after each triangulation and elimination round. We showed in Lemma 3.15 that after the l th triangulation round all rank l nodes have disjoint varieties, and thus there are at most $|\mathbf{V}(F \cap \mathbb{K}[X_l])| \leq q^k$ of them. Consider now the l th elimination round, and let us see that all the resulting rank p nodes (x_p parent of x_l) also have disjoint varieties, and thus the same bound holds.

Assume by induction that all rank p nodes have disjoint varieties before the l th elimination round. Let F_p be a rank p node (before the elimination), and let T_1, \dots, T_k be its adjacent rank l nodes. We just need to show that the new rank p nodes $F_p \cup \text{elim}_p(T_1), \dots, F_p \cup \text{elim}_p(T_k)$ have disjoint varieties (or are the same). By assumption, each $T_i \subseteq \mathbb{K}[X_l]$ defines a maximal (or prime) ideal, and thus $\text{elim}_p(T_i) \subseteq \mathbb{K}[X_l \setminus \{x_l\}]$ also defines a maximal ideal. Therefore, $\mathbf{V}(\text{elim}_p(T_i)), \mathbf{V}(\text{elim}_p(T_j))$ are either equal or disjoint, and it follows that the same holds for $\mathbf{V}(F_p \cup \text{elim}_p(T_i)), \mathbf{V}(F_p \cup \text{elim}_p(T_j))$. ■

A.2. Proofs from section 4.

Lemma A.2. *Let \mathbb{L} be a ring, and let $f \in \mathbb{L}[y]$ be a monic univariate polynomial. Let $\phi : \mathbb{L}[y] \rightarrow \mathbb{L}[y]$ be an endomorphism such that $\phi(f) = f$ and $\deg(\phi(h)) \leq \deg(h)$ for any $h \in \mathbb{L}[y]$. Then $\phi(h \bmod f) = \phi(h) \bmod f$ for any $h \in \mathbb{L}[y]$.*

Proof. Consider the Euclidean division $h = qf + r$, where $q, r \in \mathbb{L}[y]$ and $\deg(r) < \deg(f)$. Then $\phi(h) = \phi(q)f + \phi(r)$ and $\deg(\phi(r)) \leq \deg(r) < \deg(f)$, so this is the Euclidean division of $\phi(h)$. It follows that $\phi(h \bmod f) = \phi(r) = \phi(h) \bmod f$. ■

Proof of Lemma 4.7. We proceed by induction on l . The base case, $l = 0$, is clear. Assume now that the lemma holds for some l , and let us prove it for $p := l + 1$. Let f_p be a rank p node, and let $f_{l,1}, f_{l,2}, \dots, f_{l,k}$ be its adjacent rank l nodes. Let us denote by ϕ_l the functional that plugs in the values $\hat{x}_0, \dots, \hat{x}_l$. By induction, we know that

$$H(f_{l,i}) = \phi_l \left(\sum_{C_{l,i}} r_{C_{l,i}} h \bmod C_{l,i} \right),$$

where the sum is over all $f_{l,i}$ -subchains $C_{l,i}$. Note that the algorithm sets

$$H(f_p) = \phi_p \left(\sum_i r_{l,i} H(f_{l,i}) \bmod f_p \right),$$

where ϕ_p is the functional that plugs in the value \hat{x}_p . Therefore,

$$H(f_p) = \phi_p \left(\sum_i r_{l,i} \phi_l \left(\sum_{C_{l,i}} r_{C_{l,i}} h \bmod C_{l,i} \right) \bmod f_p \right) = \phi_p \left(\phi_l \left(\sum_i \sum_{C_{l,i}} r_{l,i} r_{C_{l,i}} h \bmod C_{l,i} \right) \bmod f_p \right).$$

Since any f_p -subchain is of the form $C_p = C_{l,i} \cup \{f_p\}$ for some i , we can rewrite the above equation as

$$H(f_p) = \phi_p \left(\phi_l \left(\sum_{C_p} r_{C_p} h \bmod C_p \right) \bmod f_p \right),$$

where the sum is over all f_p -subchains C_p , and where $C'_p := C_p \setminus \{f_p\}$. To complete the proof we just need to see that ϕ_l commutes with $\bmod f_p$. This follows from Lemma A.2 by setting $y = x_p$ and $\mathbb{L} = \mathbb{K}[X \setminus \{x_p\}]$. ■

Proof of Lemma 4.10. Let x_{m_i} denote the main variable of h_i , which is one of the ranks where the algorithm is initialized. It is enough to prove the lemma for ranks l where the paths (in the elimination tree) starting from different x_{m_i} first meet. Thus, we restrict ourselves to some m_1, \dots, m_k such that their respective paths all meet at rank l . More precisely, we assume that $X_{m_i}^l \cap X_{m_j}^l = \{x_l\}$ for $i \neq j$, where $X_{m_i}^l$ denotes the path in the elimination tree connecting x_{m_i} to x_l .

By applying Lemma 4.7 to each h_i , it follows that the final value of $H(f_l)$ is given by plugging in the values of $\hat{x}_1, \hat{x}_2, \dots, \hat{x}_l$ in the polynomial

$$\sum_i \sum_{C_i} r_{C_i} h_i \bmod C_i,$$

where C_i is an f_l -subchain restricted to the path $X_{m_i}^l$. Let $C = \bigcup_i C_i$ be the f_l -subchain obtained by combining them. We want to show that the above expression is equal to

$$\sum_C r_C (h_1 + \dots + h_k) \bmod C.$$

Note now that h_i does not involve any variable in $X_{m_j}^l$ for $j \neq i$. Thus, $h_i \bmod C = h_i \bmod C_i$. Observe that $X_{m_i}^l, X_{m_j}^l$ have no common arcs since they only meet at level l , and thus $r_C = \prod_i r_{C_i}$. Denoting $h_{C_i} := h_i \bmod C_i$, the problem reduces to proving the following equality:

$$(4) \quad \sum_i \sum_{C_i} r_{C_i} h_{C_i} = \sum_C r_{C_1} r_{C_2} \cdots r_{C_k} (h_{C_1} + h_{C_2} + \cdots + h_{C_k}).$$

In order to prove (4), let us look at the right-hand side as a polynomial in variables h_{C_1}, \dots, h_{C_k} . Note that the coefficient of h_{C_1} in such a polynomial is

$$\sum_{C \supseteq C_1} r_{C_1} r_{C_2} \cdots r_{C_k} = r_{C_1} \prod_{i=2}^k \left(\sum_{C_i} r_{C_i} \right),$$

and we want to show that this expression reduces to r_{C_1} . Recall that the scalar coefficients r_{C_i} are normalized (this was the second modification made to Algorithm 3). It follows that $\sum_{C_i} r_{C_i} = 1$ for all i , and thus (4) holds. ■

A.3. Proofs from section 6.

Proof of Theorem 6.10. We have to show that chordality is preserved, the variety is preserved, and the chains in the output are regular systems. The proofs of first two statements are essentially the same as for the chordally zero-dimensional case (Lemma 3.11 and Lemma 3.13). It only remains to show that the chains of the output are regular systems. Proving that the chains are square-free is very similar, so we skip it.

Let X^l denote the subtree of the elimination tree consisting of x_l and its descendants. We say that an l -subchain is the subset of a chain given by nodes of rank i for some $x_i \in X^l$. We will show by induction on l that after the l th triangulation round every l -subchain is a regular system. The base case is clear. Assume that the result holds for all $l < p$. Let \mathfrak{T}_p be a rank p node obtained after the p th triangulation round. Let \mathfrak{C} be a p -subchain containing \mathfrak{T}_p ; we want to show that it is a regular system. It is easy to see that \mathfrak{C} is triangular and that condition (i) from Definition 6.6 is satisfied. We just need to check condition (ii).

Let $f \in \mathfrak{C}$ be a rank k polynomial; we want to show that $\text{init}(f)(\hat{x}^{k+1}) \neq 0$ for any $\hat{x}^{k+1} \in \mathbf{Z}(\text{elim}_{k+1}(\mathfrak{C}))$. First consider the case that $k \geq p$, which means that $f \in \mathfrak{T}_p$. The result follows from the fact that \mathfrak{T}_p is a regular system. Assume now that $k < p$, in which case there must be a child x_l of x_p such that $x_k \in X^l$. This means that f belongs to an l -subchain \mathfrak{C}_l , which is a subset of \mathfrak{C} . Let \mathfrak{C}'_l be the same l -subchain, but before the l th elimination round. By induction, we know that \mathfrak{C}'_l is a regular system, and thus $\text{init}(f)(\hat{x}^{k+1}) \neq 0$ for any $\hat{x}^{k+1} \in \mathbf{Z}(\text{elim}_{k+1}(\mathfrak{C}'_l))$. The result follows by noticing that $\mathbf{Z}(\text{elim}_{k+1}(\mathfrak{C})) \subseteq \mathbf{Z}(\text{elim}_{k+1}(\mathfrak{C}'_l))$. ■

REFERENCES

- [1] S. B. AKERS, *Binary decision diagrams*, IEEE Trans. Computers, 100 (1978), pp. 509–516.
- [2] S. ARNBORG, D. G. CORNEIL, AND A. PROSKUROWSKI, *Complexity of finding embeddings in a k -tree*, SIAM J. Algebraic Discrete Methods, 8 (1987), pp. 277–284, <https://doi.org/10.1137/0608024>.
- [3] P. AUBRY, D. LAZARD, AND M. M. MAZA, *On the theories of triangular sets*, J. Symbolic Comput., 28 (1999), pp. 105–124.

- [4] J. R. BLAIR AND B. PEYTON, *An introduction to chordal graphs and clique trees*, in Graph Theory and Sparse Matrix Computation, Springer, New York, 1993, pp. 1–29.
- [5] H. L. BODLAENDER AND A. M. KOSTER, *Combinatorial optimization on graphs of bounded treewidth*, Comput. J., 51 (2008), pp. 255–269.
- [6] H. L. BODLAENDER AND A. M. KOSTER, *Treewidth computations I. Upper bounds*, Inform. and Comput., 208 (2010), pp. 259–275.
- [7] R. E. BRYANT, *Symbolic boolean manipulation with ordered binary-decision diagrams*, ACM Computing Surveys (CSUR), 24 (1992), pp. 293–318.
- [8] D. CIFUENTES AND P. A. PARRILO, *An efficient tree decomposition method for permanents and mixed discriminants*, Linear Algebra Appl., 493 (2016), pp. 45–81.
- [9] D. CIFUENTES AND P. A. PARRILO, *Exploiting chordal structure in polynomial ideals: A Gröbner bases approach*, SIAM J. Discrete Math., 30 (2016), pp. 1534–1570, <https://doi.org/10.1137/151002666>.
- [10] J. A. DE LOERA, J. LEE, P. N. MALKIN, AND S. MARGULIES, *Hilbert’s Nullstellensatz and an algorithm for proving combinatorial infeasibility*, in Proceedings of the 21st International Symposium on Symbolic and Algebraic Computation, ISSAC’08, ACM, New York, 2008, pp. 197–206.
- [11] R. DECHTER, *Constraint Processing*, Morgan Kaufmann, San Francisco, CA, 2003.
- [12] W. DECKER, G. GREUEL, G. PFISTER, AND H. SCHÖNEMANN, SINGULAR 3-1-6—a computer algebra system for polynomial computations, <http://www.singular.uni-kl.de>, 2012.
- [13] P. DIACONIS, D. EISENBUD, AND B. STURMFELS, *Lattice walks and primary decomposition*, in Mathematical Essays in Honor of Gian-Carlo Rota, B. E. Sagan and R. P. Stanley, eds., Progr. Math. 161, Birkhäuser Boston, Boston, MA, 1998, pp. 173–193.
- [14] S. N. EVANS, B. STURMFELS, AND C. UHLER, *Commuting birth-and-death processes*, Ann. Appl. Probab., 20 (2010), pp. 238–266.
- [15] S. GAO, *Counting Zeros over Finite Fields Using Gröbner Bases*, Master’s thesis, Carnegie Mellon University, Pittsburgh, PA, 2009.
- [16] J. HERZOG, T. HIBI, F. HREINSDÓTTIR, T. KAHLE, AND J. RAUH, *Binomial edge ideals and conditional independence statements*, Adv. Appl. Math., 45 (2010), pp. 317–333.
- [17] E. HUBERT, *Notes on Triangular Sets and Triangulation-Decomposition Algorithms I: Polynomial Systems*, Springer, New York, 2003.
- [18] T. KAHLE, *Decompositions of binomial ideals*, Ann. Inst. Statist. Math., 62 (2010), pp. 727–745.
- [19] M. KALKBRENER, *A generalized Euclidean algorithm for computing triangular representations of algebraic varieties*, J. Symbolic Comput., 15 (1993), pp. 143–167.
- [20] E. KALTOFEN, *Polynomial factorization 1987–1991*, in Proceedings of the 1st Latin American Symposium on Theoretical Informatics, I. Simon, ed., Springer, Berlin, Heidelberg, 1992, pp. 294–313.
- [21] D. E. KNUTH, *Combinatorial Algorithms, Part 1*, Art Comput. Program. 4A, Addison–Wesley, Boston, MA, 2011.
- [22] S. R. KOSARAJU, *Decidability of reachability in vector addition systems*, in Proceedings of the 14th Annual ACM Symposium on Theory of Computing, STOC ’82, ACM, New York, 1982, pp. 267–281.
- [23] S. L. LAURITZEN AND D. J. SPIEGELHALTER, *Local computations with probabilities on graphical structures and their application to expert systems*, J. Roy. Statist. Soc. Ser. B, 50 (1988), pp. 157–224.
- [24] D. LAZARD, *Solving zero-dimensional algebraic systems*, J. Symbolic Comput., 13 (1992), pp. 117–131.
- [25] F. LEMAIRE, M. M. MAZA, AND Y. XIE, *The RegularChains library*, in Maple Conference (Waterloo, Canada), Vol. 5, 2005, pp. 355–368.
- [26] M. M. MAZA, *On triangular decompositions of algebraic varieties*, presented at the MEGA-2000 Conference, NAG, UK, 2000.
- [27] M. MONAGAN AND R. PEARCE, *Sparse polynomial division using a heap*, J. Symbolic Comput., 46 (2011), pp. 807–822.
- [28] D. J. ROSE, R. E. TARJAN, AND G. S. LUEKER, *Algorithmic aspects of vertex elimination on graphs*, SIAM J. Comput., 5 (1976), pp. 266–283, <https://doi.org/10.1137/0205021>.
- [29] F. ROULLIER, *Solving zero-dimensional systems through the rational univariate representation*, Appl. Algebra Engrg. Comm. Comput., 9 (1999), pp. 433–461.
- [30] W. STEIN ET AL., *Sage: Open Source Mathematical Software*, <http://www.sagemath.org>, 2008.
- [31] B. STURMFELS, *Solving Systems of Polynomial Equations*, CBMS Regional Conf. Ser. Math. 97, AMS, Providence, RI, 2002.

- [32] L. VANDENBERGHE AND M. S. ANDERSEN, *Chordal graphs and semidefinite optimization*, Found. Trends Optim., 1 (2014), pp. 241–433.
- [33] R. VILLARREAL, *Monomial Algebras*, 2nd ed., CRC Press, Boca Raton, FL, 2015.
- [34] J. VON ZUR GATHEN AND J. GERHARD, *Modern Computer Algebra*, Cambridge University Press, Cambridge, UK, 2013.
- [35] D. WANG, *Computing triangular systems and regular systems*, J. Symbolic Comput., 30 (2000), pp. 221–236.
- [36] D. WANG, *Elimination Methods*, Springer, New York, 2001.
- [37] D. WANG, *Elimination Practice: Software Tools and Applications*, Imperial College Press, London, 2003.
- [38] I. WEGENER, *Branching Programs and Binary Decision Diagrams: Theory and Applications*, Discrete Math. Appl. 4, SIAM, Philadelphia, 2000.