

2008


Dot Product Representations of Graphs

Gregory Minton
Harvey Mudd College

Recommended Citation

Minton, Gregory, "Dot Product Representations of Graphs" (2008). *HMC Senior Theses*. 209.
https://scholarship.claremont.edu/hmc_theses/209

This Open Access Senior Thesis is brought to you for free and open access by the HMC Student Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in HMC Senior Theses by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.



Dot Product Representations of Graphs

Gregory Minton

Kimberly Tucker, Advisor

Michael E. Orrison, Reader

May, 2008

HARVEY MUDD
COLLEGE

Department of Mathematics

Copyright © 2008 Gregory Minton.

The author grants Harvey Mudd College the nonexclusive right to make this work available for noncommercial, educational purposes, provided that this copyright statement appears on the reproduced materials and notice is given that the copying is by permission of the author. To disseminate otherwise or to republish requires written permission from the author.

Abstract

We introduce the concept of dot product representations of graphs, giving some motivations as well as surveying the previously known results. We extend these representations to more general fields, looking at the complex numbers, rational numbers, and finite fields. Finally, we study the behavior of dot product representations in field extensions.

Contents

Abstract	iii
Acknowledgments	vii
1 Introduction	1
1.1 Graph Representations	1
1.2 Modeling Friendships	2
1.3 The Inverse Problem	3
2 Background	5
2.1 Definitions	5
2.2 Representations over Arbitrary Fields	8
2.3 Real Representations	12
3 Representations over \mathbb{Q} and \mathbb{C}	15
3.1 Exact Representations	15
3.2 Asymptotic Representations	16
4 Finite Fields	21
4.1 Matrix Characterizations	21
4.2 Singularity and Subadditivity	29
4.3 The Field of Order Two	42
5 Field Extensions	49
5.1 Algebraic Extensions	50
5.2 Example Graphs	54
6 Conclusions	59
Bibliography	61

Acknowledgments

I would first like to thank my advisor, Dr. Kimberly Tucker, for her advice and support throughout this research project. I also want to thank my second reader, Dr. Michael Orrison, for his comments and suggestions. Finally, the entire Harvey Mudd College Mathematics Department deserves thanks for the untold little things behind every course offered.

Chapter 1


Introduction

1.1 Graph Representations

While the definition of a graph is a collection of vertices together with a set of edges among them, it is often more convenient to describe a particular graph some other way. For example, the vertex and edge sets

$$V = \{v_1, v_2, v_3\} \quad \text{and} \quad E = \{v_1 \sim v_2, v_1 \sim v_3, v_2 \sim v_3\}$$

define a particular graph $G = (V, E)$. However, if our goal is to describe G in a way that humans can quickly understand, a picture would probably be preferable:



If our goal is simply to define G in as little space as possible, then $G = K_3$ (notation which is understood to mean the complete graph on 3 vertices) would be our preferred representation.

In computer science, often the goal is to store a graph in such a way that finding all of the neighbors of a given vertex is computationally fast. This suggests the adjacency list representation, where a graph is stored as a set of lists (one for each vertex) of adjacent vertices. In our running example,

Vertex	Adjacent Vertices
v_1	v_2, v_3
v_2	v_1, v_3
v_3	v_1, v_2

Another representation encountered in computer science is the adjacency matrix: we store an array with a row and column for each vertex, where the value in the i^{th} row and j^{th} column is 1 if the i^{th} and j^{th} vertices are adjacent, and 0 otherwise.

$$G = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}.$$

This representation is well suited if we want to quickly answer questions of whether two given vertices are adjacent.

Clearly, the way we choose to represent a particular graph depends heavily on what our goal is. A representation for a graph must completely specify the entire graph, and generically we would like to use a representation which (1) allows us to quickly recover the information we want about the graph, and (2) is as small as possible. A *dot product representation* of a graph is one particular method for representing a graph which is motivated by social modeling.

1.2 Modeling Friendships

As a toy example, imagine that three people are polled for their feelings on pets, sports, and ice cream. We use a scale where 0 denotes apathy and larger numbers signify a stronger (positive) feeling on the subject. The responses for our three people, denoted persons A, B, and C, are indicated in the following table.

	A	B	C
pets?	1	-1	0
sports?	1	1	0.5
ice cream?	0	1	0.5

If we assume that these three topics fully describe a person's preferences, and that friends tend to have similar preferences, then this table should be enough to predict the likelihood of friendships between A, B, and C. In fact, if we treat a person's preferences as a vector in \mathbb{R}^3 , then one way of mathematically encoding the statement "person P's preferences agree with person Q's preferences" is "person P's preference vector has a large projection onto person Q's preference vector." With this association in mind, we postulate the following method for predicting friendships:

If persons P and Q have preference vectors v_P and v_Q , then the probability that they will be friends is given by $v_P \cdot v_Q$.

In our example above, this specifies that A and B will not be friends, B and C will be friends, and there is a 50% chance that A and C will be friends. This loosely agrees with our intuition. A and B will fight over pets as much as they bond over sports, so (one could think) have no net basis for a friendship. B and C have similar tastes in sports and ice cream, which supports their predicted friendship. Finally, the only topic on which neither A nor C is apathetic is sports. C does not feel very strongly about the issue, so this may or may not lead to a friendship.

Work presented in Nickel (2007) shows that, given randomly generated preference vectors¹, the resulting social network of friendships follows the clustering and power law behavior expected of a social network. This result supports the intuitive basic plausibility of our model, and justifies research into preference vectors as a way of modeling social networks.

1.3 The Inverse Problem

Since preference vectors predict a certain social network, from a graph representation point of view it is natural to consider the inverse problem:

Given a certain social network (generically, a graph), what preference vectors would produce that network?

It is worth emphasizing that our prescription for using preference vectors to build a graph is non-deterministic (we actually define a probability for each edge), but to use preference vectors as a representation we must be able to reconstruct the same graph every time. Thus, in the inverse problem we are looking for preference vectors such that a pair of vectors has dot product 1 between adjacent vertices, and dot product 0 between non-adjacent vertices. This will be precisely the definition of a dot product representation of a graph, when we formally define it in the following chapter.

Having discovered a type of graph representation which is well suited for social networks, we recall that one of the goals of a representation is to be as small as possible. The size of a graph representation in preference vectors is determined by the number of preferences we have to list (i.e. the

¹The vectors are generated from a particular distribution which guarantees inner products are always between 0 and 1.

dimension of the vectors). Thus, the fundamental quantity which drives our study of dot product representations is the *minimum dimension*. This quantity may also be interpreted as a measure of the “social complexity”, as it is the smallest number of distinct preferences that we would need to poll from a population to determine all friendships. It is natural to think of a social network which could be described entirely by preferences on one issue as simpler than one which requires knowing preferences on ten issues.

Dot product representations are studied in Tucker and Scheinerman (2006a) and Tucker and Scheinerman (2006b), and their theory is developed in depth in Tucker (2007). All of these looked at representations where the representation vectors come from a vector space over the real numbers, as this is the natural field to consider when discussing preferences — while it makes sense to say that my preference for sports is 0.31, it does not make sense to say that my preference for sports is $1 + 2i$. However, we always look to generalize mathematical concepts, and thus this thesis is concerned with looking at dot product representations over different fields.

Chapter 2 states the major definitions and gives a survey of the results known for representations over \mathbb{R} . In Chapter 3, we consider representations over \mathbb{C} and \mathbb{Q} . Chapter 4 studies finite fields, where the bulk of our new results are found. In Chapter 5 we investigate behavior under field extensions. Finally, in Chapter 6 we give conclusions and avenues for future research.

Before closing the introduction, we note that dot product representations are not the only representations involving vectors. Lovász and Vesztegombi (1999) gives a general survey of so-called geometric representations. The geometric representations which are closest to dot product representations are orthogonal representations, which are studied in Lovász et al. (1989). Somewhat analogous to this work, orthogonal representations over finite fields are studied in Peeters (1996).

Chapter 2

Background

In this chapter we present the basic definitions in the study of dot product representations. We also state the major results over \mathbb{R} , as well as the known results which trivially extend to arbitrary fields.

2.1 Definitions

We first note for the record the formal definition of a graph. We always consider simple (undirected, with no self-loops) graphs.

Definition 2.1. *A graph G is a finite set of vertices V and a set of edges E , where each edge in E is an unordered pair of distinct vertices in V .*

For a graph $G = (V, E)$ and two vertices $v_1, v_2 \in V$, we say that v_1 and v_2 are adjacent (and write $v_1 \sim v_2$) if the pair $\{v_1, v_2\}$ is in E . We will often refer to the adjacency matrix for a graph G , which was mentioned in Section 1.1. To introduce notation, the adjacency matrix A for a graph G on vertices $\{v_1, \dots, v_n\}$ is the $n \times n$ matrix with entry in the i^{th} row and j^{th} column

$$A_{i,j} = \begin{cases} 1 & \text{if } v_i \sim v_j \\ 0 & \text{otherwise.} \end{cases}$$

As we outlined in Section 1.3, the definition of a dot product representation is as follows.

Definition 2.2. *Given a graph G on vertices $V = \{v_1, \dots, v_n\}$ and a field \mathbb{F} , a(n) (exact) dot product representation X of G , of dimension d over \mathbb{F} , is a mapping $X : V \rightarrow \mathbb{F}^d$ such that for $i \neq j$, $X(v_i) \cdot X(v_j) = 1$ if $v_i \sim v_j$, and $X(v_i) \cdot X(v_j) = 0$ otherwise.*

A few remarks are in order. The real-valued case, as studied in the previous literature, corresponds to $\mathbb{F} = \mathbb{R}$; such a representation will be called a real dot product representation. (Similarly, a representation over any field \mathbb{F} is an \mathbb{F} -valued dot product representation.) Whenever not specified, “representation” always refers to a dot product representation.

Our definition uses a dot product on the vector space \mathbb{F}^d . In the real or rational case, the standard dot product is used. When working over the complex numbers, we will often want to use the Hermitian inner product

$$(z_1, \dots, z_d) \cdot (w_1, \dots, w_d) = \sum_{i=1}^d z_i^* w_i.$$

Representations over \mathbb{C} using this inner product are called Hermitian representations. In finite fields, we define the dot product to be the same sum as in the real case; however, this is *not* an inner product on the vector space. For example, nonzero vectors can have zero dot product with themselves.

While a representation is defined as a map from the set of vertices to a vector space \mathbb{F}^d , we may (and often do) instead view it as a matrix in $M_{d,n}(\mathbb{F})$ whose columns are given by the representation vectors¹. We use X for both constructs, and which object we are discussing should always be clear from context. For example, the following are two ways of writing the same representation for K_3 :

$$X : \{v_1, v_2, v_3\} \rightarrow \mathbb{R}^2, \quad X(v_1) = X(v_2) = \begin{pmatrix} 3/5 \\ 4/5 \end{pmatrix}, X(v_3) = \begin{pmatrix} 5/3 \\ 0 \end{pmatrix}$$

and

$$X = \begin{pmatrix} 3/5 & 3/5 & 5/3 \\ 4/5 & 4/5 & 0 \end{pmatrix}.$$

As the above definition is called an exact representation, the terminology anticipates another sort of representation. Indeed, while an exact representation requires that the dot products be *exactly* one or zero, in certain fields we can define the concept of a representation in which dot products get arbitrarily close to their desired value.

Definition 2.3. *Given a graph G on vertices $V = \{v_1, \dots, v_n\}$ and a field \mathbb{F} , an asymptotic dot product representation of G , of dimension d over \mathbb{F} , is a sequence of mappings $\{X_k\}$ where each $X_k : V \rightarrow \mathbb{F}^d$ such that for $i \neq j$,*

$$\lim_{k \rightarrow \infty} (X_k(v_i) \cdot X_k(v_j))$$

¹As with the adjacency matrix, this association requires fixing an order on the set of vertices. Such an order is usually implied from the vertex labels.

equals 1 if $v_i \sim v_j$, and equals 0 otherwise.

Again, a few remarks are necessary for this definition. While a natural metric exists on real, rational, and complex valued vector spaces (and thus asymptotic representations are well-defined objects), it is less clear what a limit means in finite fields. Regardless of the metric we choose, since a finite field has finitely many elements, a sequence has a limit iff it becomes constant. Thus, asymptotic dot product representations are *not* an interesting concept over finite fields. We will study them over \mathbb{R} , \mathbb{Q} , and \mathbb{C} only; over the complex numbers, we will always use the Hermitian inner product.

Since there are finitely many pairs of vertices, convergence “pointwise” of each dot product implies convergence “uniformly” of all the dot products; thus if $\{X_k\}$ is an asymptotic dot product representation, then given $\epsilon > 0$ there exists a positive integer N such that

$$\left| X_n(v_i) \cdot X_n(v_j) - \lim_{k \rightarrow \infty} (X_k(v_i) \cdot X_k(v_j)) \right| < \epsilon$$

for all $n \geq N$ and all pairs of distinct vertices v_i, v_j .

As mentioned in the introduction, the primary quantity of interest is the minimum dimension of a dot product representation.

Definition 2.4. For a field \mathbb{F} and a graph G , $\text{dp}_{\mathbb{F}}(G)$ [$\text{dp}_{\mathbb{F}}^*(G)$] is the minimum dimension d such that \mathbb{F}^d admits an exact [asymptotic] dot product representation of G .

The quantity $\text{dp}_{\mathbb{F}}(G)$ is pronounced “the dot product dimension of G over \mathbb{F} .”

We close this section with some terminology used repeatedly in this thesis. Throughout,

- G is a graph,
- A is the adjacency matrix for the graph G ,
- n is the number of vertices in G ,
- D is a diagonal matrix,
- \mathbb{F} is a field, and \mathbb{F}_{p^e} is the finite field of order p^e .

Recall that there exists a unique finite field for each order p^e , where p is prime and $e \geq 0$. For a summary of the theory of finite fields, see Dummit and Foote (1999).

2.2 Representations over Arbitrary Fields

Most of the results of this section may be found in Tucker (2007) with the same proof. However, in that work the field was limited to \mathbb{R} , so we restate the proofs in our present, more general, context.

First, the parallel definition of a representation as a matrix gives us a simple alternate characterization.

Proposition 2.5. *The matrix X is an exact dot product representation of G iff $X^T X = A + D$, where A is the adjacency matrix of G and D is some diagonal matrix.*

Proof. Let the columns of X be c_1, \dots, c_n . From basic properties of matrix multiplication, $X^T X$ is the matrix whose i, j entry is $c_i \cdot c_j$. Thus, X is a representation iff the i, j entry ($i \neq j$) of $X^T X$ is 1 if the i^{th} and j^{th} vertices are adjacent, and 0 otherwise; this is precisely the same as saying $X^T X$ equals A , except possibly on the diagonal. \square

If X^\dagger denotes the conjugate transpose of the complex-valued matrix X , then by the same argument X is a Hermitian representation of G iff $X^\dagger X = A + D$.

Recall that an orthogonal transformation on a vector space is defined as a linear transformation U for which $(Uv) \cdot (Uw) = v \cdot w$ for all vectors v, w . (When using the Hermitian inner product on \mathbb{C} , our use of the term orthogonal is equivalent to unitary.) As a matrix, an equivalent definition is that U is orthogonal if $U^T U$ is the identity. Orthogonal transformations do not affect representations, as the following two lemmas show.

Lemma 2.6. *If X is an exact dot product representation into \mathbb{F}^d for a graph G and U is an orthogonal transformation on \mathbb{F}^d , then UX is also an exact dot product representation of G .*

Proof. Since X is a representation, by Proposition 2.5, $X^T X = A + D$ for some diagonal matrix D . Now

$$(UX)^T (UX) = X^T U^T U X = X^T (U^T U) X = X^T X = A + D,$$

so appealing again to Proposition 2.5, UX is also a representation. \square

Lemma 2.7. *Given an asymptotic representation $\{X_1, X_2, \dots\}$ of a graph G , any infinite subsequence is also an asymptotic representation of G . Further, for any sequence $\{U_1, U_2, \dots\}$ of orthogonal transformations, $\{U_1 X_1, U_2 X_2, \dots\}$ is also an asymptotic representation of G .*

Proof. The first statement is a standard property of limits of sequences; the second follows by the same argument as Lemma 2.6 since orthogonal transformations do not affect dot products. \square

We now present results about the dot product dimension which hold in general over all fields. For the class of paths, where the path P_n on n vertices is the graph on vertices $\{v_1, \dots, v_n\}$ with $v_i \sim v_{i+1}$, $i = 1, 2, \dots, n-1$, the dot product dimension is independent of the field we work over.

Proposition 2.8. *Let P_n be the path on n vertices. For any field \mathbb{F} , $\text{dp}_{\mathbb{F}}(P_n) = n - 1$.*

Proof. We first demonstrate an $(n - 1)$ -dimensional representation explicitly. Associate with each vertex a vector

$$X(v_j) = \begin{cases} e_1 & \text{if } j = 1 \\ e_{j-1} + e_j & \text{if } 1 < j < n \\ e_{n-1} & \text{if } j = n \end{cases} ,$$

where $\{e_i\}$ is the standard basis on \mathbb{F}^{n-1} . Dot products may be checked easily, confirming that this is a representation for P_n . This shows $\text{dp}_{\mathbb{F}}(P_n) \leq n - 1$.

Now consider any representation for P_n . We inductively prove that $\{X(v_1), \dots, X(v_k)\}$ is a linearly independent set for $k = 1, 2, \dots, n - 1$. The base case of $k = 1$ is immediate since v_1 has neighbors, and thus cannot have a zero vector. For the inductive step, assume $\{X(v_1), \dots, X(v_{k-1})\}$ is linearly independent for $1 < k \leq n - 1$. Notice that none of v_1, \dots, v_{k-1} are adjacent to v_{k+1} . Thus, for any coefficients $\alpha_1, \dots, \alpha_k$,

$$\left(\sum_{j=1}^k \alpha_j X(v_j) \right) \cdot X(v_{k+1}) = \sum_{j=1}^k \alpha_j (X(v_j) \cdot X(v_{k+1})) = \alpha_k,$$

since $X(v_j) \cdot X(v_{k+1}) = 0$ unless $j = k$, when it equals 1. Thus, if the linear combination of $X(v_j)$ is going to be zero, we must have $\alpha_k = 0$; however, by assumption the only linear combination of $\{X(v_1), \dots, X(v_{k-1})\}$ which gives zero is trivial, so the combined set $\{X(v_1), \dots, X(v_k)\}$ is linearly independent, as desired.

This inductive argument shows that $\{X(v_1), \dots, X(v_{n-1})\}$ is a set of $n - 1$ linearly independent vectors; therefore the dimension must be at least $n - 1$. We have $\text{dp}_{\mathbb{F}}(P_n) \geq n - 1$, completing the proof. \square

While we can give the exact dot product dimension for the path, for arbitrary graphs we can only give bounds on the dimension. The next several results present known bounds.

Proposition 2.9. *For any field \mathbb{F} , $\text{dp}_{\mathbb{F}}(G) \leq m$, where m is the number of edges in G .*

Proof. Let the vertices in G be $V = \{v_1, \dots, v_n\}$. Associate each coordinate in the vector space \mathbb{F}^m with one of the edges in G . Define a mapping $X : V \rightarrow \mathbb{F}^m$ where $X(v_i)$ has a 1 in the coordinate for each edge incident upon v_i , and 0 in each other. Then checking dot products, X is indeed a representation for G (of dimension m), so the minimum dimension $\text{dp}_{\mathbb{F}}(G)$ of any representation is at most m . \square

Lemma 2.10. *For a given field \mathbb{F} and positive integer $n \geq 1$, let $\{p_{i,j}\}_{1 \leq i < j \leq n}$ be an arbitrary set of elements of \mathbb{F} . There exists a set $\{X_1, \dots, X_n\}$ of vectors in \mathbb{F}^{n-1} such that $X_i \cdot X_j = p_{i,j}$ for all $i < j$.*

Proof. We construct the set of vectors iteratively. Let the standard basis on \mathbb{F}^{n-1} be $\{e_1, \dots, e_{n-1}\}$. We begin by defining $X_1 = e_1$; we now define X_j for $j > 1$ in terms of the previous X_i , such that $X_i \cdot X_j$ is correct for each $i < j$. Further, for any $j < n$ we will construct X_j with zero entries after position j and a 1 in position j ; this will be an important property to maintain for our proof.

Denote by \hat{X} the truncation of vector X to the first $j - 1$ positions. By our invariant property, all positions after $j - 1$ are zero for each vector X_1, \dots, X_{j-1} ; thus $X_r \cdot X_s = \hat{X}_r \cdot \hat{X}_s$ for each $r < s < j$. Construct a $(j - 1) \times (j - 1)$ matrix M with i^{th} row given by \hat{X}_i and let b be the $(j - 1)$ -dimensional vector $b = (p_{1,j}, p_{2,j}, \dots, p_{j-1,j})$. Again applying the invariant property, M will be a (lower) triangular matrix where the diagonal entries are all 1; thus M is invertible (with determinant 1), so there exists a (unique) vector x such that $Mx = b$. Let X be the extension of x to a $(n - 1)$ -dimensional vector by padding the extra coordinates with zero, and define $X_j = X + e_j$ (or $X_j = X$ if $j = n$).

For this definition of X_j , notice for any $i < j$

$$X_i \cdot X_j = \hat{X}_i \cdot \hat{X}_j = \hat{X}_i \cdot x = \text{row } i \text{ of } Mx = \text{row } i \text{ of } b = p_{i,j}$$

as desired.

We may always construct the next X_j , so there does exist a set $\{X_1, \dots, X_n\}$ with the desired dot products. \square

As an immediate corollary of this construction, we get the following general bound (the best possible in full generality, as Proposition 2.8 shows) on the dot product dimension.

Corollary 2.11. *For any field \mathbb{F} and graph G on n vertices, $\text{dp}_{\mathbb{F}}(G) \leq n - 1$.*

Proof. Let $V = \{v_1, \dots, v_n\}$ be the vertex set of graph G . Using Lemma 2.10, let $\{X_1, \dots, X_n\}$ be a set of $(n - 1)$ -dimensional vectors with dot products $X_i \cdot X_j$ equal to 1 if $v_i \sim v_j$ and zero otherwise. Then $X : V \rightarrow \mathbb{F}^{n-1}$ with $X(v_i) = X_i$ defines an $(n - 1)$ -dimensional dot product representation of G , showing $\text{dp}_{\mathbb{F}}(G) \leq n - 1$. \square

The path is an example of a graph where all n vertices are “different”, and in this case the best bound we can place on the dot product dimension is $n - 1$. However, if some of the vertices are “the same”, we may expect to be able to do better. The following proposition demonstrates this.

Definition 2.12. *Given a graph G , two vertices are strong twins if (1) they are adjacent, and (2) they have the same set of adjacent vertices.*

Proposition 2.13. *Let G be a graph which is not complete, and assume S is a set of strong twins. Then for any field \mathbb{F} , $\text{dp}_{\mathbb{F}}(G) \leq n - |S|$.*

Proof. For a vertex $v \notin S$, we use the notation $v \sim S$ if $v \sim s$ for some vertex $s \in S$ (and thus, by the definition of strong twins, for all $s \in S$). Note that the assumption that G is not complete guarantees that $G - S$ is not empty. Applying Lemma 2.10, find an $(n - |S| - 1)$ -dimensional mapping X on the $n - |S|$ vertices in $G - S$ such that

$$X(v) \cdot X(w) = \begin{cases} 1 & \text{if } v \sim w \\ 0 & \text{otherwise} \end{cases} - \begin{cases} 1 & \text{if } v \sim S \text{ and } w \sim S \\ 0 & \text{otherwise} \end{cases}$$

for all vertices $v \neq w$ in $G - S$.

Now extend the dimension of each $X(v)$ by one; more explicitly, let $\hat{X}(v)$ be an $(n - |S|)$ -dimensional vector, the lower $n - |S| - 1$ entries of which are $X(v)$. Make the first entry of $\hat{X}(v)$ equal to 1 if $v \sim S$, and 0 otherwise; finally, define the mapping $\hat{X}(s) = (1, 0, 0, \dots)$ for each $s \in S$.

We claim that \hat{X} is an $(n - |S|)$ -dimensional representation of G ; showing this just amounts to checking dot products, which is straightforward. \square

2.3 Real Representations

While the previous section recorded results in the literature which extend (effectively for free) to arbitrary fields, there are also many results whose proofs hold only over the real numbers. Proofs for these may be found in Tucker (2007), so here we simply give the statements without results.

These results are repeated here to allow later comparison with the corresponding results over different fields.

Proposition 2.14. *The cycle C_n on n vertices (i.e. the path P_n with the added edge $v_1 \sim v_n$) has $\text{dp}_{\mathbb{R}}(C_n) = n - 1$ if n is even and $\text{dp}_{\mathbb{R}}(C_n) = n - 2$ otherwise.*

In fact, the n even case of this proposition is subsumed by the following theorem, which is (loosely speaking) the most general class of graphs for which the dot product dimension is known exactly.

Theorem 2.15. *If G is a connected, bipartite graph, then $\text{dp}_{\mathbb{R}}(G) = n - 1$.*

The same bound holds for trees, as the following lemma and theorem (an immediate corollary) show.

Lemma 2.16. *If G is a graph containing a leaf ℓ , and \hat{G} is formed from G by removing ℓ , then $\text{dp}_{\mathbb{R}}(\hat{G}) = \text{dp}_{\mathbb{R}}(G) - 1$.*

Theorem 2.17. *If G is a tree on n vertices, then $\text{dp}_{\mathbb{R}}(G) = n - 1$.*

If we interpret the dot product dimension as a measure of complexity, then it makes sense that the union of two graphs should have the sum of the dimensions of the two subgraphs. Over \mathbb{R} , this is true.

Theorem 2.18. *If G is composed of disjoint subgraphs H_1 and H_2 , then $\text{dp}_{\mathbb{R}}(G) = \text{dp}_{\mathbb{R}}(H_1) + \text{dp}_{\mathbb{R}}(H_2)$.*

As motivation for our next result in this section, consider the representation $X(v_1) = (1, 1)$, $X(v_2) = (0.5, 0.5)$ for the single edge P_2 . These two representation vectors are not linearly independent; they both lie on the same line. If we rotate this line to the coordinate axis (applying an orthogonal transformation), then we get the representation $\tilde{X}(v_1) = (\sqrt{2}, 0)$, $\tilde{X}(v_2) = (0.5\sqrt{2}, 0)$. Clearly, we can remove the last coordinate, finding a representation of smaller dimension.

In fact, since there is a rotation in \mathbb{R}^2 taking any vector to a (multiple of a) coordinate basis vector, this technique always allows us to find a representation of smaller dimension whenever the representation vectors do not span the full space. The following theorem states the consequence of this.

Theorem 2.19. *If X is a real representation of G of dimension $\text{dp}_{\mathbb{R}}(G)$, then X has rank $\text{dp}_{\mathbb{R}}(G)$ (i.e. the representation vectors span the full space $\mathbb{R}^{\text{dp}_{\mathbb{R}}(G)}$).*

As a final result in this section, we give an alternate characterization for the dot product dimension of a graph over the real numbers. This characterization provides a template for which we will find similar results in the case of finite fields.

Definition 2.20. *A real, square matrix M is positive semidefinite if $v^T M v \geq 0$ for all vectors v .*

Recall that for real symmetric matrices, positive semidefiniteness is equivalent to all of the eigenvalues being nonnegative.

Theorem 2.21. *For a graph G with adjacency matrix A , $\text{dp}_{\mathbb{R}}(G)$ is the minimum rank of $A + D$ over all diagonal matrices D such that $A + D$ is positive semidefinite.*

Theorem 2.21 is a consequence of the spectral theorem for real, symmetric matrices (every real symmetric matrix may be orthogonally diagonalized over the reals). For an overview of positive semidefinite matrices and the real spectral theorem, see Horn and Johnson (1990).

Chapter 3

Representations over \mathbb{Q} and \mathbb{C}

In this chapter we look at exact and asymptotic representations over the complex numbers \mathbb{C} and the rational numbers \mathbb{Q} , which are two of the most natural fields (besides the real numbers) to consider.

3.1 Exact Representations

The first relationship we observe is an obvious bound on the dot product dimension in a field extension.

Lemma 3.1. *If the field \mathbb{E} contains the field \mathbb{F} , then for any graph G , $\text{dp}_{\mathbb{E}}(G) \leq \text{dp}_{\mathbb{F}}(G)$.*

Proof. If X is any \mathbb{F} -valued representation of G , then X is also an \mathbb{E} -valued representation. Thus the minimum dimension of any representation over \mathbb{E} is at most the minimum dimension over \mathbb{F} . \square

From Lemma 3.1, we immediately get the relationship $\text{dp}_{\mathbb{Q}}(G) \geq \text{dp}_{\mathbb{R}}(G) \geq \text{dp}_{\mathbb{C}}(G)$ for any graph G . However, as we noted earlier, over \mathbb{C} we are interested primarily in Hermitian representations instead of dot product representations. This leads us to the following definition.

Definition 3.2. *For a graph G , $\underline{\text{dp}}^{\dagger}(G)$ is the minimum dimension d such that \mathbb{C}^d admits an exact Hermitian representation of G .*

Since a real-valued representation is also a Hermitian representation, we get the inequality $\text{dp}_{\mathbb{R}}(G) \geq \underline{\text{dp}}^{\dagger}(G)$ for all graphs G . Surprisingly, we actually get equality!

Theorem 3.3. *For any graph G , $\text{dp}^\dagger(G) = \text{dp}_{\mathbb{R}}(G)$.*

Proof. Since we already have the opposite inequality, we just need to show $\text{dp}_{\mathbb{R}}(G) \leq \text{dp}^\dagger(G)$.

Let X be a complex representation of dimension $\text{dp}^\dagger(G)$. Recall that M^\dagger denotes the conjugate transpose of a matrix M . Now $X^\dagger X = A + D$ for some diagonal matrix D . If the vertex set of G is $\{v_1, \dots, v_n\}$, then the diagonal matrix D has j^{th} entry $X(v_j) \cdot X(v_j) = \|X(v_j)\|^2 \geq 0$; in particular, each entry of D is real, so $A + D$ is real.

For any real valued vector v ,

$$v^T(A + D)v = v^T X^\dagger X v = (Xv)^\dagger (Xv) = \|Xv\|^2 \geq 0,$$

so the matrix $A + D$ is positive semidefinite. Then applying Theorem 2.21,

$$\text{dp}_{\mathbb{R}}(G) \leq \text{rank}(A + D) = \text{rank}(X^\dagger X) \leq \text{rank } X \leq \text{dp}^\dagger(G),$$

as desired. \square

Unfortunately, the relationship between rational representations and complex representations is not as simple, as the following result of Tucker (2007) indicates.

Proposition 3.4. *The wheel graph W_6 has $\text{dp}_{\mathbb{R}}(W_6) = 3$ but $\text{dp}_{\mathbb{Q}}(W_6) > 3$.*

Proposition 3.4 shows that the dot product dimension over the rationals may be strictly larger than the dimension over the reals. At the moment, not much more is known about this relationship; in fact, it is not even known whether $\text{dp}_{\mathbb{Q}}(W_6) = 4$.

3.2 Asymptotic Representations

Several of the following results are developed (with roughly equivalent language and methods) in Tucker and Scheinerman (2006a) as a characterization of asymptotic representations. However, it can be leveraged to allow us to also characterize asymptotic representations over \mathbb{C} , ending up with a result paralleling Theorem 3.3.

We shall always use the Hermitian inner product for complex asymptotic representations. Thus, for notational convenience in comparing asymptotic and exact dimensions, in this section alone $\text{dp}_{\mathbb{C}}$ actually denotes dp^\dagger .

Lemma 3.5. *Assume \mathbb{F} is a complete field and let $\{X_1, \dots\}$ be an asymptotic representation of G with dimension d over \mathbb{F} . Either there exists a vertex v such that $\{X_k(v)\}_{k=1}^\infty$ is unbounded, or G admits an exact representation of dimension d over \mathbb{F} .*

Proof. Assume the first option is not true; we must prove G admits an exact representation of dimension d . View each X_k as an element in the complete space \mathbb{F}^{dn} . Since each vector $X_k(v)$ comprising a column of X_k is bounded, the total matrix is bounded; we thus have a bounded sequence in a complete metric space, so the sequence has a convergent subsequence. By Lemma 2.7, this subsequence $\{X_{k_j}\}$ is also an asymptotic representation of G . Applying this and the continuity of the dot product, for each pair of distinct vertices v_1, v_2 ,

$$\left[\lim_{j \rightarrow \infty} X_{k_j}(v_1) \right] \cdot \left[\lim_{j \rightarrow \infty} X_{k_j}(v_2) \right] = \lim_{j \rightarrow \infty} X_{k_j}(v_1) \cdot X_{k_j}(v_2) = \begin{cases} 1 & \text{if } v_1 \sim v_2 \\ 0 & \text{otherwise} \end{cases},$$

so $X = \lim_{j \rightarrow \infty} X_{k_j}$ is an exact representation of G with dimension d . \square

Lemma 3.6. *Let \mathbb{F} be a field of characteristic zero. If G is a graph containing vertex v , and $G - v$ denotes the graph formed by removing v from G , then $\text{dp}_{\mathbb{F}}^*(G - v) \leq \text{dp}_{\mathbb{F}}^*(G) \leq \text{dp}_{\mathbb{F}}^*(G - v) + 1$.*

Proof. The left inequality is immediate, since any representation of G is also a representation of $G - v$. For the right inequality, let $\{X_1, \dots\}$ be an asymptotic representation of $G - v$ of dimension d . For vertices v_1, v_2 in G , define χ_{v_1, v_2} to be 1 if $v_1 \sim v_2$ and 0 otherwise. We extend each X_k to a $(d + 1)$ -dimensional representation \widetilde{X}_k of G by extending $X_k(w)$ to

$$\widetilde{X}_k(w) = \begin{pmatrix} X_k(w) \\ \chi_{w, v}/k \end{pmatrix}$$

and letting

$$\widetilde{X}_k(v) = \begin{pmatrix} \vec{0} \\ k \end{pmatrix}.$$

For any two distinct vertices $w_1, w_2 \neq v$ we have $\widetilde{X}_k(w_1) \cdot \widetilde{X}_k(w_2) = X_k(w_1) \cdot X_k(w_2) + O(k^{-2})$. The last term vanishes as $k \rightarrow \infty$, and the first term has the appropriate limit of χ_{w_1, w_2} . For any vertex $w \neq v$, $\widetilde{X}_k(w) \cdot \widetilde{X}_k(v) = \chi_{w, v}$. Since all of the dot products are asymptotically correct, $\{\widetilde{X}_k\}$ is indeed an asymptotic representation of G ; this proves the right inequality. \square

The next result uses the notation $V(G)$ for the set of vertices in G , and $G - S$ (where $S \subset V(G)$) for G with all the vertices in S removed.

Proposition 3.7. *Let \mathbb{F} be \mathbb{R} or \mathbb{C} . Then*

$$\text{dp}_{\mathbb{F}}^*(G) = \min_{S \subset V(G)} (|S| + \text{dp}_{\mathbb{F}}(G - S)).$$

Proof. For any $S \subset V(G)$, repeated application of Lemma 3.6 says

$$\text{dp}_{\mathbb{F}}^*(G) \leq \text{dp}_{\mathbb{F}}^*(G - S) + |S| \leq \text{dp}_{\mathbb{F}}(G - S) + |S|,$$

which proves that $\text{dp}_{\mathbb{F}}^*(G)$ is at most the stated minimum value.

Now consider an asymptotic representation $\{X_k\}$ of G of dimension $\text{dp}_{\mathbb{F}}^*(G)$. Apply the following procedure:

- (1) If $\{X_k\}$ is bounded, quit.
- (2) Locate a vertex v such that $\{X_k(v)\}$ is unbounded.
- (3) Pass to a subsequence of $\{X_k\}$ such that $\{\|X_k(v)\|\}$ is an increasing (to infinity) sequence.
- (4) Apply an orthogonal transformation U_k to each X_k so that $X_k(v)$ is a positive multiple of the basis vector e_1 .
- (5) Remove the first row from each (transformed) matrix X_k , and remove the vertex v from consideration.

We claim that at each step of the above procedure, $\{X_k\}$ is an asymptotic representation of the (diminishing) graph under consideration. Step (2) is immediate from the condition of step (1); steps (3) and (4) are justified by Lemma 2.7. We only need to justify step (5). Pick k and let $X_k(v) = N_k e_1$. If this k is such that the representation is within ϵ of being exact, then (defining the χ indicator as in the previous proof) the first entry of $X_k(w)$ for any $w \neq v$ must be within ϵ / N_k of $\chi_{w,v} / N_k$. In particular, the contribution of the first entry to any dot product of two vertices other than v is $O(N_k^{-2})$. Since $N_k = \|X_k(v)\|$ increases without bound, this vanishes as $k \rightarrow \infty$. Thus, we can remove this coordinate without affecting the limiting dot products, which proves that the representation after step (5) is still valid.

Let S be the set of vertices removed during the above procedure. Since we remove a dimension each time we remove a vertex, the procedure terminates with a $(\text{dp}_{\mathbb{F}}^*(G) - |S|)$ -dimensional representation of $G - S$. Our termination condition is that the representation be bounded; by Lemma 3.5,

$G - S$ must therefore admit an exact representation of dimension $\text{dp}_{\mathbb{F}}^*(G) - |S|$. Thus,

$$\text{dp}_{\mathbb{F}}^*(G) - |S| \geq \text{dp}_{\mathbb{F}}(G - S) \implies \text{dp}_{\mathbb{F}}^*(G) \geq |S| + \text{dp}_{\mathbb{F}}(G - S),$$

which completes the proof of the proposition. \square

This proposition has as a corollary a characterization theorem for asymptotic representations over the real or complex numbers. For any graph G , there exists an asymptotic representation where the restriction to a subset of the coordinates is an exact representation of some induced subgraph $G - S$, and the remaining coordinates may be matched up one-to-one with the vertices in S . In the coordinate associated with some vertex v , $X_k(v)$ has entry k and each vector $X_k(w)$ for $w \neq v$ has entry $\chi_{v,w}/k$. For simplicity, we simply state this fact instead of writing it formally.

We do, however, explicitly state a different corollary.

Theorem 3.8. *For any graph G , $\text{dp}_{\mathbb{C}}^*(G) = \text{dp}_{\mathbb{R}}^*(G)$.*

Proof. From Proposition 3.7 and Theorem 3.3,

$$\text{dp}_{\mathbb{C}}^*(G) = \min_{S \subset V(G)} (|S| + \text{dp}_{\mathbb{C}}(G - S)) = \min_{S \subset V(G)} (|S| + \text{dp}_{\mathbb{R}}(G - S)) = \text{dp}_{\mathbb{R}}^*(G).$$

\square

We now turn to the study of rational asymptotic representations. Since the reals are the closure of the rational numbers, intuitively we expect asymptotic properties of rationals and reals to be identical. This intuition does indeed hold in the case of asymptotic dot product representations.

Theorem 3.9. *For any graph G , $\text{dp}_{\mathbb{Q}}^*(G) = \text{dp}_{\mathbb{R}}^*(G)$.*

Proof. Since $\mathbb{Q} \subset \mathbb{R}$, clearly $\text{dp}_{\mathbb{R}}^*(G) \leq \text{dp}_{\mathbb{Q}}^*(G)$. To prove the opposite direction, let $\{X_k\}$ be any asymptotic representation over \mathbb{R} , with some dimension d . For a given $k > 0$, define

$$\epsilon_k = \min \left\{ \frac{1}{k \cdot \max_v \|X_k(v)\|}, \frac{1}{k} \right\}.$$

Now define a matrix \tilde{X}_k by approximating each real number in X_k by a rational with an error of at most ϵ_k ; this is possible since \mathbb{Q} is dense in \mathbb{R} .

Consider a given dot product $\tilde{X}_k(v) \cdot \tilde{X}_k(w)$. It is composed of a summation of d terms, each of which is of the form

$$(a + \delta a)(b + \delta b) = ab + b(\delta a) + a(\delta b) + (\delta a)(\delta b),$$

where a and b were the corresponding entries in $X_k(v)$ and $X_k(w)$, and the error terms δa and δb are at most ϵ_k in magnitude. Because $\|X_k(v)\|$ is an upper bound on each entry in the vector $X_k(v)$, the two middle terms are each bounded above by $\frac{1}{k}$. Further, the final term is bounded above by $\frac{1}{k^2}$. Thus,

$$|\tilde{X}_k(v) \cdot \tilde{X}_k(w) - X_k(v) \cdot X_k(w)| \leq d \left(\frac{2}{k} + \frac{1}{k^2} \right) \rightarrow 0,$$

so $\{\tilde{X}_k\}$ forms a rational asymptotic representation of G , with the same dimension d . This shows $\text{dp}_{\mathbb{Q}}^*(G) \leq \text{dp}_{\mathbb{R}}^*(G)$, and completes the proof. \square

Thus, over the complex, real, and rational numbers (the three fields for which we study asymptotic representations in this thesis), the asymptotic dimension is identical. Further, it is completely determined by the exact dot product dimensions of induced subgraphs.

Chapter 4

Finite Fields

In this chapter we look at dot product representations over finite fields. We find that there are two different cases to consider: finite fields with characteristic two, and finite fields with characteristic larger than two. In particular, we give a matrix characterization of the dot product dimension in both cases. We then discuss some behavior exhibited by dot product representations over finite fields which was not seen in the real case, and close with an in-depth study of the simplest field, \mathbb{F}_2 .

4.1 Matrix Characterizations

We first seek a characterization, similar to Theorem 2.21, for the dot product dimension of a graph over a finite field. Recall that the characterization in the real case is that $\text{dp}_{\mathbb{R}}(G)$ is the minimum rank of $A + D$ over all diagonal D such that $A + D$ is positive semidefinite. We thus want to identify some analogue to a matrix being positive semidefinite over a finite field, and expect to find a characterization for the dot product dimension which involves this analogue.

While there are many properties of positive semidefinite, symmetric real matrices, the one which is really used in the proof of the real characterization (see Tucker (2007)) is that positive semidefinite matrices M may be factored as $M = L^T L$ for some matrix L . The literature on matrices over finite fields also discusses such factorizations. The following two theorems are from MacWilliams (1969).

Theorem 4.1. *If M is a symmetric, invertible matrix over a finite field of characteristic two, then M may be factored as $M = L^T L$ for a square matrix L iff M has at least one nonzero term on the main diagonal.*

Theorem 4.2. *If M is a symmetric, invertible matrix over a finite field of characteristic larger than two, then M may be factored as $M = L^T L$ for a square matrix L iff $\det A$ is a square in the field.*

Examining these theorems gives us some intuition about matrix characterizations for finite fields: (1) the cases of characteristic two and characteristic larger than two are likely to be different, and (2) the characteristic two case is probably simpler. Both of these hypotheses are, in fact, true.

To use Theorems 4.1 and 4.2 to study representations, we need to be able to handle non-invertible matrices. The following proposition will be a workhorse in our study.

Proposition 4.3. *Let \mathbb{F} be a field and let M be a symmetric $n \times n$ matrix over \mathbb{F} . If M has rank $k > 0$, then there exists an $n \times k$ matrix T and a symmetric, invertible $k \times k$ matrix U such that $M = TUT^T$.*

Proof. Let the row vectors of M be $\{r_i\}_{i=1}^n$. Since M is symmetric, these are also the column vectors; thus $\text{img } A$, which is the span of the columns, is $\text{span}\{r_i\}$. Let $\{b_j\}_{j=1}^k$ be a basis for $\text{span}\{r_i\}$, and observe in particular that k is the correct size for $\{b_j\}$.

Construct the $n \times k$ matrix T by using the b_j as column vectors; i.e. $T = \begin{pmatrix} b_1 & \dots & b_k \end{pmatrix}$. Let $v \in \mathbb{F}^n$ be some vector; then

$$\begin{aligned} v \text{ in null space of } T^T &\iff T^T v = 0 \\ &\iff b_j \cdot v = 0 \text{ for each } j \\ &\iff r_i \cdot v = 0 \text{ for each } i \\ &\iff Mv = 0 \\ &\iff v \text{ in null space of } M. \end{aligned}$$

Note that the step between $b_j \cdot v$ and $r_i \cdot v$ is true because $\text{span}\{b_j\} = \text{span}\{r_i\}$.

By the rank/nullity theorem, $\ker M$ has dimension $n - k$, and thus since M and T^T have the same kernel, $\dim \text{img } T^T = n - (n - k) = k$.

Now we let $\{f_j\}_{j=1}^k$ be a linearly independent set of vectors in \mathbb{F}^n which is also linearly independent with $\ker M$; thus a basis for $\ker M$ together with $\{f_j\}$ gives a full basis on \mathbb{F}^n . T^T is an onto map, so evaluating T^T on a basis for \mathbb{F}^n must give a spanning set for \mathbb{F}^k . However, evaluating on the basis just discussed ($\{f_j\}$ together with a basis on $\ker M$) gives $\{T^T f_1, T^T f_2, \dots, T^T f_k, 0, 0, \dots, 0\}$, since $\ker T^T = \ker M$. Thus, $\{T^T f_j\}_{j=1}^k$ is a spanning set for \mathbb{F}^k ; by checking the size, it must be a basis.

Now the image of T is the span of its columns, which is (as discussed above) the image of M . Further, by the rank/nullity theorem, $\dim \ker T + \dim \text{img } T = k$, and since the image has dimension k , T must have a trivial kernel. Thus T has a (unique, linear) left inverse $T^{-1} : \text{img } M \rightarrow \mathbb{F}^k$.

We finally have all the pieces we need. Define U as the matrix associated with the following linear transformation on \mathbb{F}^k :

$$T^T f_j \mapsto T^{-1}(Mf_j) \text{ for } j = 1, 2, \dots, k.$$

We have specified the image of a basis for \mathbb{F}^k , so we do have a well-defined linear map with a unique associated matrix. Now for each f_j ,

$$(TUT^T)f_j = TU(T^T f_j) = T(T^{-1}(Mf_j)) = Mf_j,$$

and for any vector $v \in \ker M$, $(TUT^T)v = TU(T^T v) = TU(0) = 0 = Mv$. Thus, TUT^T and M agree on a basis for \mathbb{F}^n , so are equal maps.

The only remaining claims of the proposition are that U is both invertible and symmetric. Since U is $k \times k$, if it is not invertible then it has rank $< k$; however, then $M = TUT^T$ must have rank $< k$, which contradicts the definition of k . For symmetry, since we know M is symmetric we immediately find

$$TUT^T = M = M^T = TU^T T^T.$$

Recall T has a left inverse T^{-1} , the transpose of which immediately gives a right inverse for T^T . Multiplying the above equation on either side by these inverses, we find $U = U^T$, as desired. \square

We now consider the two cases separately.

4.1.1 Characteristic Two

We will need a few preliminary results in addition to Theorem 4.1 to give our characterization of the dot product dimension. First is a result which is given in MacWilliams (1969) as a lemma for proving Theorem 4.1; we will need it here as it helps make sense of the nonzero diagonal condition. Since the proof is simple, we reproduce it as well.

Proposition 4.4. *Let M be a symmetric matrix over the finite field \mathbb{F} of order 2^m , $m \geq 1$. There exists a vector v such that $v^T M v \neq 0$ iff M has a nonzero entry on its diagonal.*

Proof. We first prove (\Leftarrow) . Let the entries in M be $a_{i,j}$ and the columns c_1, \dots, c_n , and assume $a_{k,k} \neq 0$. Then choosing $v = e_k$,

$$v^T M v = v^T c_k = a_{k,k} \neq 0.$$

We now prove (\Rightarrow) by contrapositive. Choose an arbitrary (column) vector $v = (v_1, \dots, v_n)$ and assume M has all zeros on the diagonal. Then

$$\begin{aligned} v^T M v &= (v_1, \dots, v_n) \begin{pmatrix} & \vdots & \\ \cdots & a_{i,j} & \cdots \\ & \vdots & \end{pmatrix} \begin{pmatrix} v_1 \\ \vdots \\ v_n \end{pmatrix} \\ &= \sum_{i=1}^n \sum_{j=1}^n a_{i,j} v_i v_j \\ &= 2 \sum_{i < j} a_{i,j} v_i v_j + \sum_{i=1}^n a_{i,i} v_i^2 \\ &= 0, \end{aligned}$$

where we used symmetry of A , the assumption $a_{i,i} = 0$ for all i , and the fact that \mathbb{F} has characteristic 2. This completes the proof. \square

As another preparatory lemma, we give a result which states that the condition of a nonzero diagonal can always be satisfied for a small penalty in rank.

Lemma 4.5. *Let A be an $n \times n$ matrix with rank k over some field. There exists a nonzero diagonal matrix D such that $A + D$ has rank at most $k + 1$.*

Proof. Notice that the statement is trivial if $k = n - 1$ or $k = n$, so assume $k < n - 1$.

Let N be the null space of A ; then

$$N = \{n \text{ s.t. } An = 0\}$$

is a vector space of dimension $n - k$ by the rank/nullity theorem. Let $\{n_1, \dots, n_{n-k}\}$ be a basis for N . Further, we assume without loss of generality that for any $j < n - k$, n_j has zero entry in the n^{th} position; i.e. $n_j \cdot e_n = 0$. To justify this, select any vector with nonzero entry and reindex it to be n_{n-k} ; then subtract off the appropriate multiple of this vector from each other to get a new basis that satisfies the property.

Now let D be the diagonal matrix that is all zeros except for a 1 in the bottom-right position. Then for each $j < n - k$,

$$(A + D)n_j = An_j + Dn_j = 0 + (n_j \cdot e_n)e_n = 0,$$

so $A + D$ has at least $n - k - 1$ linearly independent vectors in its kernel, and thus has rank at most $k + 1$. \square

We are now in a position to state and prove a characterization of the dot product dimension over finite fields of characteristic two. As expected, it follows the rough structure of Theorem 2.21, using the minimum rank of $A + D$ over all diagonal D such that $A + D$ is “positive semidefinite.”

Theorem 4.6. *Let \mathbb{F} be a finite field of characteristic two, and G any graph which has at least one edge. Then, if A is the adjacency matrix of G , $\text{dp}_{\mathbb{F}}(G)$ is the minimum rank of $A + D$ over all nonzero diagonal matrices D .*

Proof. Note that the assumption of G having an edge means $A + D$ will never have rank zero for any diagonal matrix D , which is important in the details below.

We first prove $\text{dp}_{\mathbb{F}}(G) \leq$ the minimum rank. Let D be a nonzero diagonal matrix such that $A + D$ has rank k ; then by Proposition 4.3, there exists an $n \times k$ matrix T and a symmetric, invertible U such that $A + D = TUT^T$. Since the diagonal of $A + D$ is given by D , which is nonzero, there exists a vector v with $v^T(A + D)v \neq 0$ by Proposition 4.4. Then $w = T^T v$ must satisfy $w^T U w = (v^T T)U(T^T v) = v^T(A + D)v \neq 0$, so again by Proposition 4.4, U must not have all zeros on the diagonal. This implies that U may be factored as $U = MM^T$ for some $k \times k$ matrix M . The rows of TM give a k -dimensional representation of G , so $\text{dp}_{\mathbb{F}}(G)$ is indeed \leq the minimum rank of $A + D$ over all nonzero D .

Conversely, let X be a representation of G of the minimum dimension $d = \text{dp}_{\mathbb{F}}(G)$, viewed as an $n \times d$ matrix with the representation vectors in the rows. Since G has an edge, note that $d > 0$. Then $XX^T = A + D$ for some diagonal matrix D . Let the rank of $A + D$ be k and notice $k \leq d$ since X has rank at most d . If D is always nonzero, then the minimum rank of $A + D$ over nonzero D is $\leq k \leq \text{dp}_{\mathbb{F}}(G)$, as desired.

Thus we only have the case remaining that $D = 0$. Applying Proposition 4.3, since A is symmetric, there exists an $n \times k$ matrix T and a symmetric, invertible $k \times k$ matrix U such that $XX^T = A = TUT^T$.

Further, since (by the construction in Proposition 4.3) T has linearly independent columns, there is a $k \times n$ left inverse T^{-1} such that $T^{-1}T = \text{id}$.

Notice that $(T^{-1})^T$ is a right inverse for T^T . Multiplying both sides of the above equation by these matrices,

$$T^{-1}XX^T(T^{-1})^T = (T^{-1}X)(T^{-1}X)^T = U.$$

We have factored U into the product of a $k \times d$ matrix and its transpose. However, in the same vein as previous arguments, since A has no zeros on its diagonal, $v^T Av = 0$ for all vectors v by Proposition 4.4. Now

$$T^{-1}A(T^{-1})^T = U,$$

so for any k -dimensional vector w ,

$$w^T U w = w^T T^{-1} A (T^{-1})^T w = ((T^{-1})^T w)^T A ((T^{-1})^T w) = 0.$$

Thus, U has all zeros on its diagonal by Proposition 4.4, so may *not* be factored as the product of a $k \times k$ matrix with its transpose. Since we have a factorization as a $k \times d$ with its transpose, we must have $k \neq d \implies k < d$.

Since A has rank k , by the previous lemma there exists a nonzero diagonal D such that $A + D$ has rank $\leq k + 1 \leq d$. Thus $d = \text{dp}_{\mathbb{F}}(G)$ is \geq the smallest rank of any $A + D$ with nonzero diagonal, completing the proof. \square

Theorem 4.6 gives us an algorithm for computing the dot product dimension over a finite field of characteristic two: we simply loop through all of the (nonzero) diagonal matrices D , and take the minimum rank of $A + D$. The rank of a matrix can be computed quickly, but there are exponentially many diagonal matrices to try. Thus this algorithm takes running time exponential in the number of vertices. However, for relatively small graphs the running time is still manageable; for example, a graph on 7 vertices over \mathbb{F}_2 only requires finding the rank of 127 matrices. While tedious for a human, this can be done by a computer very quickly.

4.1.2 Characteristic Larger than Two

In contrast with the characteristic two case, where the analogue to positive semidefiniteness is just having nonzero diagonal, the remaining case is not as simple.

Knowing that we will use Proposition 4.3 and Theorem 4.2 in our characterization, we are led to the following definition.

Definition 4.7. Let \mathbb{F} be a finite field with characteristic larger than 2. For an $n \times n$ symmetric matrix M over \mathbb{F} with rank r , let T be an $n \times r$ matrix and U an $r \times r$ symmetric, invertible matrix such that $M = TUT^T$. Define the character of M , denoted $\chi_{\mathbb{F}}(M)$, to be 1 if $\det U$ is a square and zero otherwise. We let $\chi_{\mathbb{F}}(0) = 1$. When the field \mathbb{F} is understood, we write $\chi(M) = \chi_{\mathbb{F}}(M)$.

This definition is a priori dependent on the choice of factorization TUT^T ; thus, we must first show that the character is well-defined.

Proposition 4.8. For any symmetric matrix M , $\chi(M)$ is well-defined.

Proof. Assume M is not zero, and let $M = T_1U_1T_1^T = T_2U_2T_2^T$ be any two factorizations with the T_i matrices of dimension $n \times r$ and the U_i symmetric, invertible matrices of dimension $r \times r$. Recall that each T_i must have columns providing a basis for the column space of M ; thus the columns of T_1 are spanned by the columns of T_2 , so there exists some $r \times r$ matrix K_1 such that $T_1 = T_2K_1$. Similarly, the columns of T_2 are spanned by the columns of T_1 , so there exists an $r \times r$ matrix K_2 such that $T_2 = T_1K_2$. Finally, recall that T_1 and T_2 have left inverses, which we denote E_1 and E_2 , respectively. Thus

$$\begin{aligned} T_1U_1T_1^T &= T_2U_2T_2^T \\ T_2K_1U_1K_1^TT_2^T &= T_2U_2T_2^T \\ E_2(T_2K_1U_1K_1^TT_2^T)E_2^T &= E_2(T_2U_2T_2^T)E_2^T \\ K_1U_1K_1^T &= U_2 \end{aligned}$$

so $\det U_2 = (\det K_1)^2 \det U_1$. Thus $\det U_1$ is a square if and only if $\det U_2$ is a square, so $\chi(M)$ is indeed well-defined. \square

If we take the matrices with $\chi = 1$ as the ‘‘positive semidefinite’’ matrices in this situation, then the following characterization theorem suggests itself.

Theorem 4.9. Let \mathbb{F} be a finite field with characteristic $p > 2$. Then for any graph G with adjacency matrix A , $\text{dp}_{\mathbb{F}}(G)$ is the minimum rank of $A + D$ over all diagonal matrices D such that $\chi(A + D) = 1$.

Proof. For simplicity, define d to be the minimum rank of $A + D$ over all diagonal D such that $\chi(A + D) = 1$.

For any diagonal D , we may write $A + D = TUT^T$, where the columns of T form a basis for the column space of $A + D$ and U is a symmetric,

invertible matrix. Recall that symmetric, invertible matrices U may be decomposed as $U = MM^T$ for an invertible square matrix M if and only if $\det U$ is a square. Thus, if $\chi(A + D) = 1$ (i.e. $\det U$ is a square), then we may write $A + D = TUT^T = (TM)(TM)^T$, so the rows of TM give a representation of dimension $\text{rank}(A + D)$. Thus, $\text{dp}_{\mathbb{F}}(G) \leq d$.

To prove the theorem we need to show $\text{dp}_{\mathbb{F}}(G) \geq d$. Any representation with dimension less than d must also have rank less than d ; thus let X be any representation of G with $\text{rank } X < d$. Since X is composed of n vectors and has $\text{rank} < n$, there must be some vertex v whose vector $X(v)$ is a linear combination of the other representation vectors:

$$X(v) = \sum_{w \neq v} \alpha_w X(w).$$

To be a representation, we must have $X^T X = \hat{A}$ where \hat{A} is the sum of some diagonal matrix and the adjacency matrix A . If X^T was onto, then we could write $\hat{A} = (X^T)\text{id}(X^T)^T$ where the columns of X^T are linearly independent and the identity is invertible with a square determinant, so $d \leq \text{rank } \hat{A} \leq \text{rank } X$. We assumed this was not the case, so X^T is not onto and thus $\text{rank } X$ is strictly less than the dimension of the representation.

For a vertex w , let $r(w)$ be the corresponding row of \hat{A} . Then $r(v)$ is linearly dependent on the other rows, since

$$\begin{aligned} r(v) &= (X(v))^T (X(v_1), \dots, X(v_n)) \\ &= \left(\sum \alpha_w X(w) \right)^T (X(v_1), \dots, X(v_n)) \\ &= \sum \alpha_w r(w). \end{aligned}$$

Further, if $e_v \in \mathbb{F}^n$ is the unit basis vector in the coordinate corresponding to vertex v , then $\{r(w)\}$ cannot span e_v : if it did, then the corresponding linear combination of the representation vectors would give a vector orthogonal to all $X(w)$ for $w \neq v$ but not to $X(v)$, and this is a contradiction as $X(v)$ is linearly dependent on the $X(w)$.

Let $D = \text{diag } e_v$. Then for nonzero x , $\hat{A} + xD$ is formed from \hat{A} by (1) removing the row corresponding to v , which does not change the rank since this row was linearly dependent on the others; (2) inserting a row with the vector $x e_v$, which adds one to the rank since e_v is not in the row span; and (3) adding back the original row corresponding to v to this newly inserted row, which again does not change the rank since the original row was and is spanned by the other rows. Thus $\text{rank}(\hat{A} + xD) = \text{rank } \hat{A} + 1$.

Let $r = \text{rank } \hat{A}$ and let $\{b_1, \dots, b_{r+1}\}$ be a basis for the column space (equivalently, row space) of the matrix $\hat{A} + xD$, where $b_{r+1} = e_v$ (possible

since e_v is in the column space). Then $\{b_1, \dots, b_r\}$ is a basis for the column space of \hat{A} , so we may write

$$\hat{A} = (b_1 \ \dots \ b_r) \hat{U} \begin{pmatrix} b_1^T \\ \vdots \\ b_r^T \end{pmatrix}$$

for some symmetric, invertible matrix \hat{U} .

Let $x = \det \hat{U} \neq 0$. Then

$$(b_1 \ \dots \ b_r \ b_{r+1}) \begin{pmatrix} \hat{U} & \vec{0} \\ \vec{0}^T & x \end{pmatrix} \begin{pmatrix} b_1^T \\ \vdots \\ b_r^T \\ b_{r+1}^T \end{pmatrix} = \hat{A} + xD.$$

Thus, we have constructed a factorization $\hat{A} + xD = TUT^T$ where T has linearly independent columns and U is a symmetric, invertible matrix with determinant $(\det \hat{U})^2$, a square, so $d \leq \text{rank}(\hat{A} + xD) = \text{rank} \hat{A} + 1 \leq \text{rank} X + 1$, which in turn is \leq the dimension of the representation X , and so $d \leq \text{dp}_{\mathbb{F}}(G)$.

This completes the proof. \square

Theorem 4.9 again gives us an exponential time algorithm for computing the dot product dimension. An applet which actually uses these algorithms to compute the dot product dimension over a general finite field can be found at¹

<http://www.math.hmc.edu/~gminton/thesis/DPApplet.html>

and Figure 4.1 gives a screenshot of this applet.

4.2 Singularity and Subadditivity

In this section we explore the differences between Theorems 2.17, 2.18, and 2.19 (which were results over the real numbers) and the corresponding results over finite fields. In each case, we find that finite fields exhibit behavior more complicated than their real counterpart.

¹URL valid as of writing.




Figure 4.1: Screenshot of the dot product dimension applet

4.2.1 Singularity

From Theorem 2.19, we know that representations of minimum dimension over the real numbers must span the entire vector space. This property fails to hold over finite fields, and is at the heart of most of the odd behavior studied in this section. Intuitively, this failure is because nonzero vectors can be orthogonal to themselves. We study the concept in depth for fields of characteristic two (with a few results specific to \mathbb{F}_2), and then close with a general theorem.

To formalize the subject of discussion, we make the following definitions.

Definition 4.10. A representation $X : \{v_1, \dots, v_n\} \rightarrow \mathbb{F}^d$ is deficient if

$$\dim \text{span}\{X(v_j)\} < d.$$

Definition 4.11. The graph G is singular over a field \mathbb{F} if it contains an edge and admits a deficient representation of dimension $\text{dp}_{\mathbb{F}}(G)$.

Now consider singularity over finite fields of characteristic two. Directly from the characterization of the dot product dimension, we get the following results.

Proposition 4.12. Let \mathbb{F} be a finite field with characteristic two. If G is singular over \mathbb{F} , then $\text{rank } A = \text{dp}_{\mathbb{F}}(G) - 1$.

Proof. Recall our characterization: $\text{dp}_{\mathbb{F}}(G)$ is the minimum rank of $A + D$ over all nonzero diagonal matrices D . Further, in the proof of this result we noted that for any matrix B , there exists a nonzero diagonal matrix E such that $\text{rank}(B + E) \leq \text{rank } B + 1$. Thus

$$\text{dp}_{\mathbb{F}}(G) = \min_{D \neq 0} \text{rank}(A + D) \leq \text{rank } A + 1,$$

so $\text{rank } A \geq \text{dp}_{\mathbb{F}}(G) - 1$.

Now let X be a deficient representation of dimension $\text{dp}_{\mathbb{F}}(G)$. Then viewed as a matrix, $X^T X = A + D$ for some diagonal D . Further, $\text{rank}(A + D) \leq \text{rank } X < \text{dp}_{\mathbb{F}}(G)$. To avoid a contradiction with the definition of $\text{dp}_{\mathbb{F}}(G)$ as the minimum rank, we must have $D = 0$; then $\text{rank } A < \text{dp}_{\mathbb{F}}(G)$.

Since $\text{rank } A$ is an integer, it must equal $\text{dp}_{\mathbb{F}}(G) - 1$. \square

Corollary 4.13. *Suppose \mathbb{F} is finite with characteristic two. If X is a deficient representation of dimension $\text{dp}_{\mathbb{F}}(G)$, then X has rank $\text{dp}_{\mathbb{F}}(G) - 1$ and $A = X^T X$.*

Proof. View X as a matrix; then $X^T X = A + D$ for some diagonal matrix D . Then $\text{rank}(A + D) \leq \text{rank } X < \text{dp}_{\mathbb{F}}(G)$, where $\text{dp}_{\mathbb{F}}(G)$ is the minimum rank of $A + D$ for any nonzero diagonal matrix D . Thus $D = 0$, so $\text{rank } X \geq \text{rank } A = \text{dp}_{\mathbb{F}}(G) - 1$ by the previous proposition. Then $\text{rank } X = \text{dp}_{\mathbb{F}}(G) - 1$, as desired. \square

These conditions allow us to write down several more properties of singular graphs, which greatly restrict where we may look for examples of singularity.

Lemma 4.14. *Suppose \mathbb{F} is a finite field of characteristic two, and v is a vector over \mathbb{F} . Let $e = (1, 1, \dots, 1)$. We have $v \cdot v = 0$ iff $v \cdot e = 0$.*

Proof. Let $v = (v_1, \dots, v_n)$. Since \mathbb{F} has characteristic two, the map $a \mapsto a^2$ splits over addition; thus

$$(v \cdot e)^2 = (v_1 + \dots + v_n)^2 = v_1^2 + \dots + v_n^2 = v \cdot v,$$

which immediately proves the statement. \square

Lemma 4.15. *Suppose G is singular over a finite field \mathbb{F} of characteristic two. The dimension $\text{dp}_{\mathbb{F}}(G)$ is odd.*

Proof. Let $d = \text{dp}_{\mathbb{F}}(G)$ and let X be a deficient representation of dimension d . Label the columns of X by v_1, \dots, v_n . Then by Corollary 4.13, X has rank $d - 1$. Further, $X^T X = A$. Since A has zero entries on the diagonal, each column of X must be self-orthogonal. Thus, defining $e = (1, 1, \dots, 1)$ and using Lemma 4.14,

$$0 = v_j \cdot v_j \implies 0 = v_j \cdot e.$$

Now recall that the orthogonal complement of any given nonzero vector is a vector space of one less dimension. Thus,

$$\mathcal{O} = \{v : v \cdot e = 0\}$$

is a vector space of dimension $d - 1$ which contains each column v_j . However, the columns span a space of dimension $d - 1$, so the column space of X is \mathcal{O} .

Now notice that $A = X^T X$, and all of these matrices have rank $d - 1$ (by Proposition 4.12 and Corollary 4.13). Thus, the image of X must intersect trivially with the kernel of X^T . The kernel of X^T is the space of vectors orthogonal to each v_j , which includes e as seen above; the image of X is \mathcal{O} . If d is even, then e is self-orthogonal, so $e \in \mathcal{O}$, which provides a nontrivial intersection. Thus d must be odd, as desired. \square

Proposition 4.16. *Over a finite field of characteristic two, any representation of minimum dimension spans either the full vector space or the orthogonal complement of $(1, 1, \dots, 1)$.*

Proof. From Corollary 4.13, a representation X of minimum dimension d may be deficient only if $\text{rank } X = d - 1$ and $X^T X = A$; i.e. $X^T X$ has zeros on the diagonal. Thus, each vector $X(v_j)$ is self-orthogonal, so by Lemma 4.14

$$0 = X(v_j) \cdot (1, 1, \dots, 1)$$

for each vertex v_j . Thus the image of X is contained in the orthogonal complement of $(1, 1, \dots, 1)$. However, this orthogonal complement has the dimension $d - 1 = \text{rank } X$, so the columns of X do indeed span the orthogonal complement of $(1, 1, \dots, 1)$. \square

The following result deserves some explanation; since $\text{rank } A = \text{dp}_{\mathbb{F}}(G) - 1$, it actually states that each row of the adjacency matrix may be written as a linear combination of the other rows. Since weak twins (nonadjacent vertices with the same neighbors) have the same rows in the adjacency matrix, this heuristically suggests that sets of weak twins are likely to appear in singular graphs.

Lemma 4.17. *Let \mathbb{F} be a finite field of characteristic two. Suppose G is singular and let $\mathcal{R} = \{r_1, \dots, r_n\}$ be the rows of A . For any $j = 1, 2, \dots, n$, $\dim \text{span } \mathcal{R} \setminus \{r_j\} = \text{dp}_{\mathbb{F}}(G) - 1$.*

Proof. From Proposition 4.12, $\text{rank } A = \dim \text{span } \mathcal{R} = \text{dp}_{\mathbb{F}}(G) - 1$. Since adding a vector to a collection can only increase the dimension of its span by 0 or 1, if we assume the contradiction of the problem statement then there exists a j such that $\dim \text{span } \mathcal{R} \setminus \{r_j\} = \text{dp}_{\mathbb{F}}(G) - 2$. Let d be the vector which is zero except for a one in position j , and let $D = \text{diag } d$. Then

$$\dim \text{span}(\mathcal{R} \setminus \{r_j\}) \cup \{r_j + d\} \leq \dim \text{span}(\mathcal{R} \setminus \{r_j\}) + 1 = \text{dp}_{\mathbb{F}}(G) - 1.$$

However, $\dim \text{span}(\mathcal{R} \setminus \{r_j\}) \cup \{r_j + d\} = \text{rank}(A + D)$, and D is nonzero so $\text{rank}(A + D) \geq \text{dp}_{\mathbb{F}}(G)$. This provides the desired contradiction. \square

Our final result of this sort is a combinatorial lemma restricting the size of singular graphs.

Lemma 4.18. *There are no singular graphs over \mathbb{F}_2 with 4 or fewer vertices.*

Proof. Assume G is a singular graph with $n \leq 4$ vertices. Then $\text{dp}_{\mathbb{F}_2}(G) \leq n - 1 \leq 3$, and by Lemma 4.15 we know $\text{dp}_{\mathbb{F}_2}(G)$ is odd; thus $\text{dp}_{\mathbb{F}_2}(G) = 1$ or $\text{dp}_{\mathbb{F}_2}(G) = 3$.

If $\text{dp}_{\mathbb{F}_2}(G) = 1$, then a deficient representation of dimension $\text{dp}_{\mathbb{F}_2}(G)$ must be composed of all zero vectors. Singular graphs are required to have edges, so this cannot be a valid representation; thus $\text{dp}_{\mathbb{F}_2}(G) > 1$.

The only remaining case is $\text{dp}_{\mathbb{F}}(G) = 3$ and $n = 4$. Thus by Proposition 4.12, $\text{rank } A = 2$. If A is a 4×4 matrix with rank 2, then it has a null space of dimension 2; let n_1 and n_2 span this null space.


Assume for contradiction that there is some coordinate j in which both n_1 and n_2 are zero. Then letting D be the diagonal matrix which is zero except for a one in position j , we see

$$(A + D)n_i = An_i = 0, \quad i = 1, 2.$$

Thus $A + D$ has a null space including n_1 and n_2 ; if $\text{nullity}(A + D) \geq 2$, then $\text{rank}(A + D) \leq 4 - 2 = 2$. However, since D is nonzero, $\text{rank}(A + D) \geq \text{dp}_{\mathbb{F}_2}(G) = 3$, which gives our contradiction.

Thus in each coordinate, there are three possibilities: (1) n_1 is one but n_2 is zero, (2) n_1 is zero but n_2 is one, or (3) both n_1 and n_2 are one. Since there are four positions, there must be two positions which fall into the same category. Let D be the diagonal matrix which only has ones in these two positions; then $(A + D)n_i = An_i = 0$ for $i = 1, 2$, so we again find a contradiction.

We have contradicted each possible case, so the original assumption of a singular graph on ≤ 4 vertices must be invalid. \square

Figure 4.2: A singular graph over \mathbb{F}_2

Lemma 4.18 is in fact the strongest such statement possible, as a brute force search through all graphs on five vertices reveals that the graph in Figure 4.2 is actually singular over \mathbb{F}_2 . The rank of its adjacency matrix is 2, but it has dot product dimension 3. While there are nondeficient representations of minimum dimension, the following example of a deficient representation (of minimum dimension) shows that this graph is singular:

$$X(v_1) = \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}, \quad X(v_2) = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}, \quad X(v_3) = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix},$$

$$X(v_4) = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \quad X(v_5) = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}.$$

In particular, these representation vectors do not span $(1,1,1)$ (note that Proposition 4.16 predicted this).

Over finite fields of characteristic two, Corollary 4.13 has as a generalization the following theorem.

Theorem 4.19. *G is singular over a finite field of characteristic two iff $\text{rank } A < \text{rank}(A + D)$ for any nonzero diagonal matrix D .*

Proof. Let \mathbb{F} be a finite field of characteristic two. Suppose G is singular and let X be a deficient representation of minimum dimension. Then $X^T X = A$, so $\text{rank } A \leq \text{rank } X = \text{dp}_{\mathbb{F}}(G) - 1 = \min_{D \neq 0} \text{rank}(A + D) - 1 < \text{rank}(A + D)$ for any nonzero D .

Consider the converse. Following the proof of our characterization for $\text{dp}_{\mathbb{F}}$, for any symmetric $n \times n$ matrix M with rank r , we can find a decomposition $M = T^T U T$ for a symmetric, invertible $r \times r$ matrix U . Now $r = \text{rank } A < \text{rank}(A + D)$ for any nonzero diagonal matrix D , so $\text{rank } A < \text{dp}_{\mathbb{F}}(G)$. Write $A = T^T U T$, where U is symmetric and invertible.

The $(r + 1) \times (r + 1)$ matrix

$$\hat{U} = \begin{pmatrix} U & 0 \\ 0 & 1 \end{pmatrix}$$

will be symmetric, invertible, and does not have all zeros on the diagonal, so there exists an $(r + 1)$ -dimensional square matrix V with $\hat{U} = V^T V$. Form \bar{V} by removing the last column of V ; then $\bar{V}^T \bar{V} = U$, so $A = (\bar{V}T)^T (\bar{V}T)$ and thus the columns of $\bar{V}T$ give a representation of G with dimension $r + 1 \leq \text{dp}_{\mathbb{F}}(G)$. However, the rank of this representation is $\leq \text{rank } \bar{V} = r$, so as desired we have constructed a deficient representation of minimum dimension. \square

The corresponding result is, in fact, true for general finite fields. We present a sequence of results which concludes with a key theorem.

Proposition 4.20. *Let G be a graph with adjacency matrix A , and \mathbb{F} a finite field of characteristic larger than 2. If D is a diagonal matrix and $\text{rank}(A + D) = r$, then $\text{dp}_{\mathbb{F}}(G) \leq r + 1$. Further, if $\text{dp}_{\mathbb{F}}(G) = r + 1$, then G is singular.*

Proof. Let T be an $n \times r$ matrix whose columns provide a basis for the column space of $A + D$. Then for some symmetric, invertible matrix U , $A + D = TUT^T$. Let $x = \det U$, and form the $(r + 1) \times (r + 1)$ matrix

$$\hat{U} = \left(\begin{array}{c|c} U & \vec{0} \\ \hline \vec{0}^T & x \end{array} \right).$$

Then $\det \hat{U} = x(\det U) = (\det U)^2$ is a square, so \hat{U} may be factored as $\hat{U} = \hat{M}\hat{M}^T$ for an invertible $(r + 1)$ -dimensional matrix \hat{M} . Now if M is the $r \times (r + 1)$ matrix whose rows are the first r rows of \hat{M} , then $MM^T = U$. Thus $A + D = TUT^T = TMM^T T^T = (TM)(TM)^T$, so the rows of TM give an $(r + 1)$ -dimensional representation of G , and thus $\text{dp}_{\mathbb{F}}(G) \leq r + 1$.

Further, since the rows of \hat{M} were all linearly independent, $\text{rank } M = \text{rank } \hat{M} - 1 = (r + 1) - 1 = r$. Thus, since T has trivial kernel, $\text{rank } TM = \text{rank } M = r < r + 1$, so the representation TM is deficient. If $\text{dp}_{\mathbb{F}}(G) = r + 1$, then TM is a representation with minimum dimension; thus G is singular. \square

The following corollary is now immediately obtained with no further proof, as we already knew the result for finite fields of characteristic two.

Corollary 4.21. *Over any finite field \mathbb{F} and any graph G with adjacency matrix A ,*

$$\min_{D \text{ is diagonal}} \text{rank}(A + D) \in \{\text{dp}_{\mathbb{F}}(G), \text{dp}_{\mathbb{F}}(G) - 1\}.$$

The main result of singularity is the following theorem.

Theorem 4.22. *A graph G is singular over a finite field \mathbb{F} iff there exists a diagonal matrix D with $\text{rank}(A + D) = \text{dp}_{\mathbb{F}}(G) - 1$.*

Proof. Theorem 4.19 implies this result for finite fields of characteristic two; thus, we only consider \mathbb{F} with characteristic larger than two.

Suppose G is singular, and let X be a deficient representation of dimension $\text{dp}_{\mathbb{F}}(G)$. Then by definition of deficiency, $\text{rank } X < \text{dp}_{\mathbb{F}}(G)$. Now since X is a representation, $X^T X = A + D$ for some D , and $\text{rank}(A + D) \leq \text{rank } X \leq \text{dp}_{\mathbb{F}}(G) - 1$. Now by Proposition 4.20, $\text{rank}(A + D) \geq \text{dp}_{\mathbb{F}}(G) - 1$, so $\text{rank}(A + D) = \text{dp}_{\mathbb{F}}(G) - 1$; this proves the (\Rightarrow) direction.

If D has $\text{rank}(A + D) = \text{dp}_{\mathbb{F}}(G) - 1$, then Proposition 4.20 directly implies that G is singular. We have shown both directions. \square

Corollary 4.21 and Theorem 4.22 provide a wealth of connections. Algorithmically, if we ignore the concern that $A + D$ must be “positive semidefinite” in the characterization of the dot product dimension, and take the minimum rank over all diagonal D , then we will have an error in our estimate of the dot product dimension of at most one. Further, a graph is nonsingular iff the minimum rank over all diagonal matrices equals the dot product dimension.

Thus singularity is the difference between our problem of computing the dot product dimension of a graph, and the more generic problem of minimizing the rank of a matrix where the diagonal is left undetermined. This is an example of a matrix completion problem: we are finding the values for undetermined matrix elements which extremize some property of the matrix. An overview of such problems is given in Laurent (2001). In Barrett et al. (2006), the authors consider the minimum rank problem over \mathbb{F}_2 where the diagonal is left undetermined; except for singularity issues, this is exactly our problem. Their work mostly looks at characterizing matrices with a given minimum rank.

4.2.2 Subadditivity

Theorem 2.18 states that the dot product dimension over \mathbb{R} is additive on disjoint graphs. We do not get as strong of a result over finite fields, leading to the following definition.

Definition 4.23. A graph G composed of components H_j is *subadditive* (on $\{H_j\}$, over \mathbb{F}) if $\text{dp}_{\mathbb{F}}(G) < \sum \text{dp}_{\mathbb{F}}(H_j)$.

While subadditivity is generic to finite fields, so far the theory has only been developed to any extent over \mathbb{F}_2 . Thus, the remainder of this section gives the results known in this case; generalizations of them are extremely likely.

The following proposition gives a class of subadditive graphs.

Proposition 4.24. If G is a singular graph, then for any n , the graph G' with components G and the complete graph K_n has $\text{dp}_{\mathbb{F}_2}(G') = \text{dp}_{\mathbb{F}_2}(G)$.

Proof. Let X be a deficient representation of G of dimension $\text{dp}_{\mathbb{F}_2}(G)$; then X has rank $\text{dp}_{\mathbb{F}_2}(G) - 1$ and by Proposition 4.16 the image of X must be the orthogonal complement of $(1, 1, \dots, 1)$. Further, from Lemma 4.15, $\text{dp}_{\mathbb{F}_2}(G)$ is odd, so $(1, 1, \dots, 1)$ is not self-orthogonal. $X(v)$ is orthogonal to $o = (1, 1, \dots, 1)$ and $o \cdot o = 1$, so defining $X(k_j) = o$ for each vertex $k_j \in K_n$ gives a representation for G' ; thus $\text{dp}_{\mathbb{F}_2}(G') \leq \text{dp}_{\mathbb{F}_2}(G)$. Finally, G' contains G , so we get the opposite inequality and thus equality. \square

In particular, this proposition together with our example from the last section give the example in Figure 4.3 of a subadditive graph, as the left component has dimension 3, the right component (a single edge) has dimension 1, and the joint graph has dimension $3 < 3 + 1$.




Figure 4.3: A subadditive graph

We lose additivity of the dot product dimension precisely because of the phenomenon of singularity: minimum dimension representations do not have to be spanning. The correct analogue of additivity is the following additivity of dimensions result.

Lemma 4.25. Let G be a graph with components $\{v_i\}$ and $\{w_j\}$ and no isolated vertices. Then in a representation X of G ,

$$\dim \text{span}\{X(v_i)\} + \dim \text{span}\{X(w_j)\} = \dim \text{span}(\{X(v_i)\} \cup \{X(w_j)\}).$$

Proof. Clearly,

$$\dim \text{span}(\{X(v_i)\} \cup \{X(w_j)\}) \leq \dim \text{span}\{X(v_i)\} + \dim \text{span}\{X(w_j)\}.$$

Let $s_1 = \dim \text{span}\{X(v_i)\}$ and $s_2 = \dim \text{span}\{X(w_j)\}$; pick a subset $\{b_p\}_1^{s_1}$ of $\{v_i\}$ such that $\{X(b_p)\}$ is linearly independent, and a similar subset $\{c_q\}_1^{s_2}$ of $\{w_j\}$.

Let $\sum \beta_p X(b_p) + \sum \gamma_q X(c_q) = 0$ be an arbitrary linear combination yielding zero. For each p , pick a vertex v_i which is adjacent to b_p and dot the above expression with v_i ; this gives

$$\sum_{z: b_z \sim v_i} \beta_z X(b_z) = 0$$

and since the $X(b_z)$ are linearly independent, each $\beta_z = 0$. In particular, $\beta_p = 0$ since $b_p \sim v_i$. Similarly, each $\gamma_q = 0$, so the trivial linear combination is all that works; this shows $\dim \text{span}(\{X(v_i)\} \cup \{X(w_j)\}) \geq s_1 + s_2$ and completes the proof. \square

We can now leverage Lemma 4.25 to give connections between subadditivity and singularity.

Proposition 4.26. *Suppose G is subadditive with components H_1, H_2 . At least one of H_1, H_2 is singular.*

Proof. Assume that neither H_1, H_2 are singular; but then the rank of any representation of H_i is $\geq \text{dp}_{\mathbb{F}_2}(H_i)$. Let X be a representation of G with minimum dimension; then X induces representations X_i on H_i and Lemma 4.25 gives

$$\text{rank } X = \text{rank } X_1 + \text{rank } X_2 \geq \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2).$$

However, since X is a minimum representation, $\text{rank } X \leq \text{dp}_{\mathbb{F}_2}(G)$. We have shown $\text{dp}_{\mathbb{F}_2}(G) \geq \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2)$, which is a contradiction since G is subadditive. \square

Proposition 4.27. *Suppose G is subadditive with components H_1, H_2 . If only one of H_1, H_2 is singular, then G is not singular.*

Proof. Suppose without loss of generality that H_1 is singular, but H_2 is not. Then the rank of any representation of H_2 is at least $\text{dp}_{\mathbb{F}_2}(H_2)$, and the rank of any representation of H_1 is at least $\text{dp}_{\mathbb{F}_2}(H_1) - 1$. Applying this, for a representation X of minimum dimension,

$$\text{rank } X = \text{rank } X_1 + \text{rank } X_2 \geq \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2) - 1.$$

However, $\text{rank } X \leq \text{dp}_{\mathbb{F}_2}(G) < \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2)$, so $\text{rank } X = \text{dp}_{\mathbb{F}_2}(G) = \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2) - 1$; thus X cannot be deficient. \square

Proposition 4.28. *Suppose G has components H_1, H_2 . If both H_1, H_2 are singular, then G is singular.*

Proof. Let A_j be the adjacency matrix for H_j . Then since both H_j are singular, $\text{rank } A_j < \text{rank } A_j + D_j$ for any nonzero diagonal matrix D_j . If we order the vertices of H_1 before the vertices of H_2 , then the adjacency matrix A of G is block diagonal:

$$A = \begin{pmatrix} A_1 & 0 \\ 0 & A_2 \end{pmatrix}.$$

Thus, for any diagonal D with block components D_1 and D_2 , $\text{rank}(A + D) = \text{rank}(A_1 + D_1) + \text{rank}(A_2 + D_2)$. With our opening observation, $\text{rank}(A + D)$ is thus strictly minimized by $D = 0$, which proves by Theorem 4.19 that G is singular. \square

Proposition 4.29. *Suppose G has components H_1, H_2 . Then $\text{dp}_{\mathbb{F}_2}(G)$ is either $\text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2)$ or $\text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2) - 1$.*

Proof. First, observe that concatenating a representation for H_1 with a representation for H_2 gives a representation for G with the sum of the dimensions; thus $\text{dp}_{\mathbb{F}_2}(G) \leq \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2)$ always holds. Assume G is subadditive; i.e. $\text{dp}_{\mathbb{F}_2}(G) < \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2)$.

Let X be a representation of G with dimension $\text{dp}_{\mathbb{F}_2}(G)$, which induces representations X_j for H_j , $j = 1, 2$. Now $\text{rank } X_j \geq \text{dp}_{\mathbb{F}_2}(H_j)$ unless H_j is singular, when $\text{rank } X_j \geq \text{dp}_{\mathbb{F}_2}(H_j) - 1$; thus, in either case, by Lemma 4.25,

$$\text{dp}_{\mathbb{F}_2}(G) \geq \text{rank } X = \text{rank } X_1 + \text{rank } X_2 \geq \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2) - 2.$$

Consider the case $\text{rank } X = \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2) - 2$; however, this requires that both H_1 and H_2 are singular, so by Proposition 4.28 G will also be singular, implying $\text{dp}_{\mathbb{F}_2}(G) = \text{rank } X + 1 = \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2) - 1$.

In every case, either $\text{dp}_{\mathbb{F}_2}(G) = \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2)$ or $\text{dp}_{\mathbb{F}_2}(G) = \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2) - 1$, as desired. \square

The preceding set of propositions may be summarized into the following theorem, whose proof is given immediately by Propositions 4.26, 4.27, 4.28, and 4.29.

Theorem 4.30. *If G is subadditive with components H_1, H_2 , then at least one is singular and G is singular iff both are. Further, $\text{dp}_{\mathbb{F}_2}(G) = \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2) - 1$. If G is not subadditive, then $\text{dp}_{\mathbb{F}_2}(G) = \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2)$.*

This theorem gives a rather full characterization of subadditivity over \mathbb{F}_2 .

4.2.3 Leaves

We now consider the related, but distinct, question of what happens when a leaf is added to or removed from a graph. One generally uses the term leaf in the context of trees; however, we generally do not assume the graph is connected or acyclic. Again, while generalizations are likely in arbitrary finite fields, in this section we mostly study representations over \mathbb{F}_2 .

Recall from Lemma 2.16 that adding a leaf always increases the dot product dimension over \mathbb{R} by 1. The same is not true in finite fields. For example, the star graph on $n \geq 3$ vertices has dot product dimension 2 over \mathbb{F}_2 , independent of n (assign the center node $(1,0)$, and each leaf $(1,1)$). Thus, removing a leaf from the star on 5 vertices (for a specific example) does not change the dimension.

We have found a few interesting results related to leaf removal; the first is a “it is almost like the real case” result, which holds in general fields.

Proposition 4.31. *Let \mathbb{F} be any field. If G' is formed from G by removing a leaf, then $\text{dp}_{\mathbb{F}}(G')$ is either $\text{dp}_{\mathbb{F}}(G)$ or $\text{dp}_{\mathbb{F}}(G) - 1$.*

Proof. Since G contains G' , $\text{dp}_{\mathbb{F}}(G) \geq \text{dp}_{\mathbb{F}}(G')$. However, given any representation of G' we may add a coordinate and give each vertex of G' a value of 0 in this coordinate except for the vertex adjacent to the leaf, which gets a 1; then giving the leaf the unit basis vector in the new coordinate gives a representation for G . Thus $\text{dp}_{\mathbb{F}}(G) \leq \text{dp}_{\mathbb{F}}(G') + 1$, completing the proof. \square

Another interesting leaf-related result belongs in spirit with the lemmas from Section 4.2.1, where we restricted which graphs could be singular.

Proposition 4.32. *Any graph containing a leaf is nonsingular over \mathbb{F}_2 .*

Proof. Let G be a graph with a leaf. Let the leaf be in j^{th} position and let its (only) neighbor be in k^{th} position. Thus the j^{th} row of A is all zeros, except for a single one in position k . Let D be the diagonal matrix with a single one in the k^{th} position; then $A + D$ is the matrix formed from A by the row

operation (row k) \mapsto (row k) + (row j) and thus $\text{rank}(A + D) = \text{rank } A$. By the condition of Theorem 4.19, this implies G cannot be singular. \square

In particular, this shows that trees cannot be singular. This has as a corollary the following intriguing theorem, which exploits results from an existing theory of the “minimum rank of a graph” to connect dot product representations across different fields.

Theorem 4.33. *Let T be a tree. Then $\text{dp}_{\mathbb{F}}(T) \geq \text{dp}_{\mathbb{F}_2}(T)$ for any field \mathbb{F} .*

Proof. Define the minimum rank of a graph G with n vertices over a field \mathbb{F} to be the minimum rank of any symmetric $n \times n$ matrix over \mathbb{F} whose i, j entry ($i \neq j$) is nonzero exactly when vertices i and j are adjacent. Let X be a dot product representation of G with minimum dimension. Then $X^T X$ is a matrix satisfying the above conditions, and $\text{rank}(X^T X) \leq \text{rank } X = \text{dp}_{\mathbb{F}}(G)$, so the minimum rank of a graph can be at most the dot product dimension.

Consider the minimum rank of any nonsingular graph G over \mathbb{F}_2 . Since the only nonzero element of \mathbb{F}_2 is 1, the collection of symmetric matrices with i, j entry nonzero iff i and j are adjacent is precisely the collection $\{A + D\}$, where D ranges over all diagonal matrices. Since we assume G is nonsingular, by Theorem 4.19, $\text{rank } A \geq \min\{\text{rank}(A + D) : D \neq 0\}$. Thus the minimum rank of G is $\min\{\text{rank}(A + D)\} = \min\{\text{rank}(A + D) : D \neq 0\} = \text{dp}_{\mathbb{F}_2}(G)$.

As shown in Chenette et al. (2007), the minimum rank of a tree is field independent. Since trees are nonsingular by Proposition 4.32, for any field \mathbb{F} ,

$$\text{dp}_{\mathbb{F}_2}(T) = \min. \text{rank of } T \text{ over } \mathbb{F}_2 = \min. \text{rank of } T \text{ over } \mathbb{F} \leq \text{dp}_{\mathbb{F}}(T).$$

\square

The interplay between singularity and leaf removal yields another result; taking a singular graph and adding a leaf *anywhere* does not change the dot product dimension! The following proposition proves this.

Proposition 4.34. *Let G be a graph with a leaf ℓ , and let \hat{G} be the graph formed by removing ℓ from G . If \hat{G} is singular, then $\text{dp}_{\mathbb{F}_2}(G) = \text{dp}_{\mathbb{F}_2}(\hat{G})$.*

Proof. As is standard, let G have n vertices; then \hat{G} has $n - 1$. Since G contains \hat{G} , $\text{dp}_{\mathbb{F}_2}(G) \geq \text{dp}_{\mathbb{F}_2}(\hat{G})$. Let \hat{A} be the adjacency matrix for \hat{G} . Then $\text{rank } \hat{A} = \text{dp}_{\mathbb{F}_2}(\hat{G}) - 1$ since \hat{G} is singular. Assume without loss of generality that the leaf ℓ is in the last position of the adjacency matrix A of G , and

that it is adjacent to the vertex which is next-to-last (so in last position in \hat{A}). Let D be the $n \times n$ diagonal matrix $D = \text{diag}(0, 0, \dots, 0, 1, 1)$; then $A + D$ may be formed from \hat{A} by adding an extra row and column of zeros (which does not change the rank) and then adding the vector $(0, 0, \dots, 0, 1, 1)$ to each of the last two rows. Since we add the same vector to both, this operation can increase the rank by at most one; thus

$$\text{rank}(A + D) \leq \text{rank } \hat{A} + 1 = \text{dp}_{\mathbb{F}_2}(\hat{G}).$$

However, since D is nonzero, $\text{dp}_{\mathbb{F}_2}(G) \leq \text{rank}(A + D)$. We have shown $\text{dp}_{\mathbb{F}_2}(\hat{G}) \leq \text{dp}_{\mathbb{F}_2}(G)$ and $\text{dp}_{\mathbb{F}_2}(G) \leq \text{dp}_{\mathbb{F}_2}(\hat{G})$, so the two are equal. \square

Note that, since removing a leaf does not always change the dot product dimension, we no longer have as simple a characterization of the dot product dimension of a tree as in the real case. For example, while the path P_6 is a tree with 6 vertices and $\text{dp}_{\mathbb{F}_2}(P_6) = 5$, the graph in Figure 4.4 is a tree with 6 vertices and dot product dimension 4. Also, as noted, the star graph with $n \geq 3$ vertices (which is a tree) has dot product dimension 2 independent of n . Clearly, the behavior of trees is no longer as simple!




Figure 4.4: A tree with dot product dimension 4

4.3 The Field of Order Two

In this section, we focus our attention exclusively on the smallest field, \mathbb{F}_2 , where an alternative and more “physical” interpretation of the dot product dimension is possible. This new interpretation allows us to prove a theorem, applicable for all finite fields of characteristic two, on which graphs have the maximum possible dot product dimension $n - 1$.

We rely on the following graph operation, which is sometimes alternatively called “toggling.”

Definition 4.35. *Given a graph G and a subset $U \subseteq V$ of the vertices of G , the subgraph complement (of G , with respect to U), denoted \overline{G}_U , is the graph on vertex*

set V where $v_1 \sim v_2$ in \overline{G}_U iff either $v_1 \sim v_2$ in G and $\{v_1, v_2\} \not\subset U$ or $v_1 \not\sim v_2$ in G and $\{v_1, v_2\} \subset U$.

Put differently, a subgraph complement is exactly what it sounds like: we take the complement of an induced subgraph, and do not affect any edges outside of that subgraph. For example, Figure 4.5 shows a graph and its subgraph complement.




Figure 4.5: A graph and its subgraph complement

Notice that subgraph complementing twice with respect to the same subset restores the original graph. Further, subgraph complements are commutative. In particular, composition of complements gives an abelian group structure to

$$\langle (G \mapsto \overline{G}_U) : U \subseteq V, \sigma = \rho \text{ iff } \sigma(G) = \rho(G) \rangle^2.$$

Some of the graph operations we are familiar with can be represented by subgraph complements; a few are summarized in the following proposition. In the language below, we are intentionally vague about the vertex set of a graph; in particular, induced subgraphs are referred to as still containing all of the vertices in the original graph (just with no edges incident upon them).

Proposition 4.36. *Let G and H be graphs.*

²For two graphs G and H and two sequences of subgraph complements σ and ρ , $\sigma(G) = \rho(G)$ iff $\sigma(H) = \rho(H)$, so equality in the group is well-defined. Since every generator has order 2 and the group is abelian, every non-identity element of the group has order 2. The group elements are in one-to-one correspondence with graphs on n vertices, so this group is actually isomorphic to $(\mathbb{Z}/2\mathbb{Z})^{n(n-1)/2}$.

1. If H differs from G on exactly m edges (i.e. $|E(H) \triangle E(G)| = m$), then H may be obtained from G by m subgraph complement operations.
2. If H is an induced subgraph of G with k vertices removed, then H may be obtained from G by $2k$ subgraph complement operations.
3. If H is an induced subgraph of G obtained by removing a set of strong twins, then H may be obtained from G by 2 subgraph complement operations.

Proof. The subgraph complement with respect to a pair of vertices simply toggles the edge between them, proving (1).

For (2), suppose that a vertex $v \in V(G)$ has neighbors $\{v_1, \dots, v_d\}$. Consider first the subgraph complement with respect to $\{v, v_1, \dots, v_d\}$ and then the subgraph complement with respect to $\{v_1, \dots, v_d\}$; this does not affect any edges except those through v (and it removes all of those), so this pair of operations indeed gives the induced subgraph where v is removed. This proves (2).

For (3), notice that the operations in (2) still work with a set of strong twins. \square

Subgraph complements are interesting for our purposes because of the following theorem.

Theorem 4.37. *Let G be a graph. $\text{dp}_{\mathbb{F}_2}(G)$ is the smallest number of subgraph complement operations necessary to transform the edgeless graph into G .*

Proof. Let G have n vertices, and number them 1 through n . Let X be a $d \times n$ matrix over \mathbb{F}_2 ; we associate X with the sequence of d subgraph complement operations, where the subset associated with the i^{th} complement is the subset of columns j for which $X_{i,j} = 1$. For example,

$$\begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{pmatrix} \rightarrow \left\{ \begin{array}{l} 1. \text{ complement w.r.t. } \{1,3\} \\ 2. \text{ complement w.r.t. } \{2,3\} \\ 3. \text{ complement w.r.t. } \{1,3\} \end{array} \right\}.$$

This establishes a bijection between $d \times n$ matrices over \mathbb{F}_2 and sequences of d subgraph complements; to complete the proof we show that X is a dot production representation of G (i.e. $X^T X$ off the diagonal equals the adjacency matrix) iff the sequence of complements associated with X takes the edgeless graph to G .

Pick a pair $i \neq j$ of vertices. Expanding the definition,

$$\begin{aligned} (X^T X)_{i,j} &= (\text{column } i \text{ of } X) \cdot (\text{column } j \text{ of } X) = \sum_k X_{k,i} X_{k,j} \\ &= (\text{number of rows in which } X_{k,i} = X_{k,j} = 1) \pmod{2}. \end{aligned}$$

Now let (C_1, \dots, C_d) be the sequence of subgraph complements associated with X , and consider these operations applied to the edgeless graph. There will be an edge between i and j iff the edge was toggled an odd number of times, which in turn happens iff $\{i, j\}$ is contained in the subsets of an odd number of the complement operations. But $\{i, j\}$ is contained in the subset of complement C_k iff $X_{k,i} = X_{k,j} = 1$. Thus edge $\{i, j\}$ will be present iff $(X^T X)_{i,j} = 1$, completing the proof. \square

Corollary 4.38. *If H can be obtained from G by k subgraph complement operations, then $\left| \text{dp}_{\mathbb{F}_2}(H) - \text{dp}_{\mathbb{F}_2}(G) \right| \leq k$.*

In particular, Corollary 4.38 combined with Proposition 4.36 immediately gives us that

1. If H differs from G on exactly m edges (i.e. $|E(H) \triangle E(G)| = m$), then $\text{dp}_{\mathbb{F}_2}(G) - m \leq \text{dp}_{\mathbb{F}_2}(H) \leq \text{dp}_{\mathbb{F}_2}(G) + m$.
2. If H is an induced subgraph of G with k vertices removed, then $\text{dp}_{\mathbb{F}_2}(G) - 2k \leq \text{dp}_{\mathbb{F}_2}(H) \leq \text{dp}_{\mathbb{F}_2}(G) + 2k$.
3. If H is an induced subgraph of G obtained by removing a set of strong twins, then $\text{dp}_{\mathbb{F}_2}(G) - 2 \leq \text{dp}_{\mathbb{F}_2}(H) \leq \text{dp}_{\mathbb{F}_2}(G) + 2$.

We also immediately find $\text{dp}_{\mathbb{F}_2}(G) - 1 \leq \text{dp}_{\mathbb{F}_2}(\overline{G}) \leq \text{dp}_{\mathbb{F}_2}(G) + 1$: the dot product dimension of a graph differs by at most one from the dimension of its complement. Such an association does not hold over all fields. For example, the complete bipartite graph $K_{n,n}$ has dimension $2n - 1$ over the reals and its complement (two copies of K_n) has dimension only 2.

We can extend this result slightly when G is singular.

Proposition 4.39. *If G is singular over \mathbb{F}_2 , then $\text{dp}_{\mathbb{F}_2}(G) - 1 \leq \text{dp}_{\mathbb{F}_2}(\overline{G}) \leq \text{dp}_{\mathbb{F}_2}(G)$.*

Proof. Let $d = \text{dp}_{\mathbb{F}_2}(G)$ and let X be a deficient representation of dimension d for G . Define the vector $e = (1, 1, \dots, 1) \in \mathbb{F}_2^d$. From Proposition 4.16, each representation vector is orthogonal to e . Further, from Lemma 4.15, d is odd, so $e \cdot e = d = 1$.

Define a representation \overline{X} by $\overline{X}(v) = X(v) + e$ for each vertex v . Then, for any pair v_i, v_j of vertices,

$$\overline{X}(v_i) \cdot \overline{X}(v_j) = X(v_i) \cdot X(v_j) + e \cdot X(v_j) + X(v_i) \cdot e + e \cdot e = X(v_i) \cdot X(v_j) + 1.$$

Since $X(v_i) \cdot X(v_j) = 1$ iff $v_i \sim v_j$ in G , we have $\bar{X}(v_i) \cdot \bar{X}(v_j) = 1$ iff $v_i \sim v_j$ in \bar{G} . Thus \bar{X} gives a d -dimensional representation of \bar{G} .

This proves the inequality on the right; the inequality on the left holds for all graphs by our arguments above. \square

It is worth noting that the subgraph complement characterization does not really give us any new mathematical tools for attacking the dot product dimension. It cannot; as we see from the proof of Theorem 4.37, it is really just a rephrasing of the same concept. However, it does provide a more intuitive language for some ideas (like the bounds above).

It also provides a more graphical interpretation of the dot product dimension. For a fixed number of vertices n , consider forming a graph whose vertex set is the collection of all graphs. Edges are drawn between graphs which are related by a single subgraph complement. The resulting structure is a $(2^n - n - 1)$ -regular graph where the distance from a graph to the edgeless graph gives its dot product dimension. For example, Figure 4.6 is the diagram for $n = 3$.




Figure 4.6: The toggling diagram for $n = 3$

There are 64 graphs on $n = 4$ vertices, so the corresponding diagram is quite large. Isomorphic graphs, though, all have the same dot product dimension. Thus we can take minors, collapsing every set of isomorphic graphs into a single representative, while still keeping most of the interesting information. The reduced diagram for $n = 4$ is in Figure 4.7

Notice that in Figures 4.6 and 4.7, only the path has the maximum dot product dimension $\text{dp}_{\mathbb{F}_2} = n - 1$. This holds for all n , and is an example




Figure 4.7: The toggling diagram for $n = 4$

of a result which is proven most cleanly with the language of subgraph complementing.

Theorem 4.40. *A graph G on n vertices has $\text{dp}_{\mathbb{F}_2}(G) = n - 1$ iff it is the path.*

Proof. The theorem can be computationally verified for all graphs on 3 or fewer vertices. Further, the (\Leftarrow) direction has already been proven (Proposition 2.8), so it is just the (\Rightarrow) direction which remains. We induct on n . Suppose inductively that it is true for graphs on $n - 1$ (with $n - 1 \geq 3$) vertices and suppose for contradiction that G is a graph on n vertices which is not the path but has $\text{dp}_{\mathbb{F}_2}(G) = n - 1$.

Define the operation of “vertex excision” to be the subgraph complement with respect to the subset consisting of a given vertex and all of its neighbors. For any vertex v , let G_v denote the excision of v from G . Since each G_v is a graph on $n - 1$ vertices, if it is not the path P_{n-1} then it has $\text{dp}_{\mathbb{F}_2}(G_v) \leq n - 3$ by the inductive hypothesis. Thus, appealing to Corollary 4.38, $\text{dp}_{\mathbb{F}_2}(G) \leq \text{dp}_{\mathbb{F}_2}(G_v) + 1 \leq n - 2$. Since $\text{dp}_{\mathbb{F}_2}(G) = n - 1$, each G_v must in fact be the path.

If G is not connected, then it has components H_1 and H_2 with n_1 and n_2 vertices, respectively, and

$$\text{dp}_{\mathbb{F}_2}(G) \leq \text{dp}_{\mathbb{F}_2}(H_1) + \text{dp}_{\mathbb{F}_2}(H_2) \leq (n_1 - 1) + (n_2 - 1) = n - 2.$$

Thus G must be connected. The path and cycle are the only connected graphs in which every vertex has degree ≤ 2 ; since G is not the path by assumption and the cycle has $\text{dp}_{\mathbb{F}_2}(C_n) = n - 2$, G must have a vertex of degree ≥ 3 .

Now fix such a vertex v in G and label the remaining vertices such that G_v is the path $v_1 \leftrightarrow v_2 \leftrightarrow \cdots \leftrightarrow v_{n-1}$. Let $N \subseteq \{v_1, \dots, v_{n-1}\}$ be the collection of neighbors of v in G .

Assume $|N| < n - 1$, and let w be a vertex not adjacent to v in G . Then the excision G_w does not affect the degree of v ; since G_w contains a vertex of degree ≥ 3 , it cannot be the path.

The only remaining case is $|N| = n - 1$, so $G_v = P_{n-1}$ is actually the complement of G . Thus G has

$$\binom{n}{2} - (n - 2) = \frac{n^2}{2} - \frac{3n}{2} + 2$$

edges. Since $n - 1 \geq 3$, there is a vertex z that had degree 2 in G_v (in particular, v_2), which will have degree $n - 3$ in G . Excising this vertex takes the complement of a subgraph on $n - 2$ vertices. Further, this subgraph is the complement of the subgraph induced from P_{n-1} by removing an interior vertex, which has $n - 4$ edges. Since complementing a graph on $n - 2$ vertices with $n - 4$ edges adds

$$\binom{n-2}{2} - 2(n-4) = \frac{n^2}{2} - \frac{9n}{2} + 11$$

edges, excision by z (taking this complement again) reduces the number of edges by the same amount. Subtracting, G_w has $3n - 9$ edges. Since $3n - 9 \neq n - 2$, G_w cannot be P_{n-1} , providing our final contradiction and proving the theorem. \square

Now by Proposition 2.8, $\text{dp}_{\mathbb{F}}(P_n) = n - 1$ holds for any field \mathbb{F} . Further, $\text{dp}_{\mathbb{F}}(G) \leq \text{dp}_{\mathbb{F}_2}(G)$ for any field $\mathbb{E} \supseteq \mathbb{F}_2$ by Lemma 3.1. Thus the above theorem actually holds over any field of characteristic two. However, it does not hold in full generality over all finite fields; for example, the star graph on four vertices has dot product dimension 3 over \mathbb{F}_3 , but is certainly not the path.

As an interesting special case, the results of this section (in particular, Theorem 4.40 and Corollary 4.38) immediately imply that the complement of the path, $\overline{P_n}$, has dot product dimension $\text{dp}_{\mathbb{F}_2}(\overline{P_n}) = n - 2$.

Chapter 5

Field Extensions

For most of this thesis, we have been implicitly considering the field to be fixed; i.e. looking at $\text{dp}_{\mathbb{F}}(G)$ as a function of the graph G . In this chapter we change the focus and instead examine $\text{dp}_{\mathbb{F}}(G)$ as a function of the field \mathbb{F} .

The first result in such a study is Lemma 3.1; the dot product dimension is (weakly) decreasing under field extensions. Since the dimension is always a nonnegative integer, there must be some minimum dimension. This motivates the following definition.

Definition 5.1. *Let p be 0 or a prime. For a graph G , the limiting dimension of G over characteristic p is*

$$\hat{\text{dp}}_p(G) = \min \{ \text{dp}_{\mathbb{F}}(G) : \mathbb{F} \text{ is a field with characteristic } p \}.$$

One bound on the limiting dimension is immediately possible, exploiting the field independence result for the path.

Proposition 5.2. *If G has diameter d , then $\hat{\text{dp}}_p(G) \geq d$ for all characteristics p .*

Proof. Recall Proposition 2.8, that the path P_n has dot product dimension $n - 1$ over any field. Now any graph with diameter d contains P_{d+1} as an induced subgraph, so for any field \mathbb{F} , $\text{dp}_{\mathbb{F}}(G) \geq \text{dp}_{\mathbb{F}}(P_{d+1}) = d$. This gives the desired bound. \square

We can get another bound from results already known, but it only holds for trees.

Proposition 5.3. *If T is a tree, then $\hat{\text{dp}}_p(T) \geq \text{dp}_{\mathbb{F}_2}(T)$ for all characteristics p .*

Proof. Let \mathbb{F} be any field with characteristic p ; then by Theorem 4.33, $\text{dp}_{\mathbb{F}}(T) \geq \text{dp}_{\mathbb{F}_2}(T)$. Thus $\hat{\text{dp}}_p(T) = \min\{\text{dp}_{\mathbb{F}}(T) : \mathbb{F} \text{ has char. } p\} \geq \text{dp}_{\mathbb{F}_2}(T)$. \square

In particular, it follows that $\hat{\text{dp}}_2 = \text{dp}_{\mathbb{F}_2}$ for all trees.

5.1 Algebraic Extensions

The limiting dimension must be achieved by some field, which raises the question of “how large” of a field is needed. For a motivating example, consider the wheel graph W_6 . As we saw in Section 3.1, it has a smaller representation over the real numbers than over the rational numbers. Using the entire field \mathbb{R} , though, is excessive; in the case of W_6 , the degree-four extension $\mathbb{Q}[\phi^{1/2}]$ (ϕ denotes the golden ratio $(1 + \sqrt{5})/2$) suffices to get a representation of dimension 3. Indeed, we expect in general that some *algebraic* extension of the characteristic field should always achieve the limiting dimension.

After developing a few preparatory results, we give a theorem proving this result.

Proposition 5.4. *Let \mathbb{E} be a field extension of \mathbb{F} with finite degree. A set of vectors $\{v_j\}_1^n$ over \mathbb{F} is linearly independent iff $\{v_j\}$ is linearly independent as a set of vectors over \mathbb{E} .*

Proof. If there exist no nontrivial coefficients $\{a_j\}$ in \mathbb{E} such that $\sum a_j v_j = 0$, then clearly there exist no nontrivial coefficients over \mathbb{F} ; thus (\Leftarrow) is immediate.

We now prove (\Rightarrow) by contrapositive. Let $\sum \alpha_j v_j = 0$ be a nontrivial linear combination with coefficients $\alpha_j \in \mathbb{E}$. View \mathbb{E} as a vector space over \mathbb{F} of dimension d ; for each j , write $\alpha_j = (a_j^1, a_j^2, \dots, a_j^d)$ with each $a_j^q \in \mathbb{F}$. Then since each vector v_j has coefficients in \mathbb{F} , multiplication is pointwise; i.e.

$$0 = \sum \alpha_j v_j = \sum (a_j^1, \dots, a_j^d) \cdot v_j = \left(\sum a_j^1 v_j, \dots, \sum a_j^d v_j \right).$$

There must exist j and q such that $a_j^q \neq 0$ (otherwise each $\alpha_j = 0$). Choosing this q , $\sum a_j^q v_j$ gives a nontrivial linear combination of $\{v_j\}$ over \mathbb{F} yielding zero. \square

This has as a corollary the following result:

Corollary 5.5. *Let \mathbb{E} be a field extension of \mathbb{F} with finite degree. Consider a set of vectors $\{v, v_1, \dots, v_n\}$ over \mathbb{F} . The condition $v \in \text{span}\{v_1, \dots, v_n\}$ is independent of the field of scalars being \mathbb{F} or \mathbb{E} .*

Proof. Without loss of generality, for an appropriate choice of k , $\{v_j\}_1^k$ gives a basis for $\{v_1, \dots, v_n\}$. By the proposition, this is still linearly independent over \mathbb{E} , so this gives a basis over either field. Thus, over either \mathbb{E} or \mathbb{F} , $v \in \text{span}\{v_1, \dots, v_n\}$ iff $\{v\} \cup \{v_j\}_1^k$ is linearly dependent. The latter condition is field-independent by the proposition, so the desired condition is as well. \square

We now need a result from algebraic geometry. Recall that the affine variety $V(I)$ defined by an ideal I in the polynomial ring $K[x_1, \dots, x_n]$ is

$$V(I) = \{x = (x_1, \dots, x_n) \in K^n : f(x) = 0 \text{ for all } f \in I\}.$$

Theorem 5.6 (Hilbert's (Weak) Nullstellensatz). *Let K be an algebraically closed field and consider the multivariate polynomial ring $K[x_1, \dots, x_n]$. An ideal J of this ring is proper iff $V(J)$ is not empty.*

See, for example, Dummit and Foote (1999) for a proof of this theorem.

Consider the ideal I generated by a finite set $\{p_1, \dots, p_k\}$ of polynomials in $K[x_1, \dots, x_n]$. It can be quickly verified that $f(x) = 0$ for all $f \in I$ iff $p_1(x) = \dots = p_k(x) = 0$. Further, an ideal is proper iff it does not contain 1; thus I is proper iff 1 is not a $K[x_1, \dots, x_n]$ -linear combination of $\{p_1, \dots, p_k\}$ (i.e. a sum $f_1 p_1 + \dots + f_k p_k$, where each coefficient f_j is itself a polynomial).

Stated in a more convenient form, the weak form of Hilbert's Nullstellensatz thus says the following: a set of polynomials has a common zero (over an algebraically closed field) iff 1 is not a (polynomial-coefficient) linear combination of them.

We are interested in the following consequence of this.

Proposition 5.7. *Let \mathbb{E} be any field and let \mathbb{F} be its characteristic field. Suppose \mathcal{A} is a finite set of algebraic equations with coefficients in \mathbb{F} . If there exists a solution to \mathcal{A} over \mathbb{E} , then there exists a finite-degree extension of \mathbb{F} containing a solution to \mathcal{A} .*

Proof. We can manipulate the equations in \mathcal{A} so the right-hand sides are all zero, and thus the problem is one of finding common roots for a set of (potentially multivariate) polynomials $\{p_j\}_1^n$.

We prove the statement by contrapositive; suppose there is no finite-degree extension of \mathbb{F} which has a common root for $\{p_j\}$. Let \mathbb{A} be the field

extension of \mathbb{F} formed by adjoining all algebraic numbers over \mathbb{F} . Then, by contradiction, \mathbb{A} cannot contain a common root for $\{p_j\}$; if it did, then adjoining the algebraic numbers used in the common root would give a finite-degree extension of \mathbb{F} .

Now \mathbb{A} is algebraically closed since a polynomial whose coefficients are algebraic has algebraic roots. By Theorem 5.6, since $\{p_j\}$ has no common root over the algebraically closed field \mathbb{A} , the function 1 may be written as a polynomial-linear combination $1 = \sum f_j p_j$, where each f_j is a polynomial over \mathbb{A} . Let $\{M_i\}_1^m$ be the collection of monomials (setting the coefficients to 1) present in $\{f_j\}$, and let \mathcal{B} be the collection of all monomials present in $\{M_i p_j\}_{i,j}$. Consider the vector space over \mathbb{A} with basis \mathcal{B} . This is a finite-dimensional space; further, since each p_j has coefficients in \mathbb{F} , the representation of each polynomial $M_i p_j$ has coefficients in \mathbb{F} .

There are a finite number of algebraic coefficients in $\{f_j\}$; thus the statement $1 = \sum f_j p_j$ means $1 \in \text{span}\{M_i p_j\}_{i,j}$, where the span is taken over a finite-degree extension of \mathbb{F} . However, by Corollary 5.5, this implies $1 \in \text{span}\{M_i p_j\}_{i,j}$ where the span is taken over \mathbb{F} . Thus we may write $1 = \sum g_j p_j$, where each g_j is a polynomial with coefficients in \mathbb{F} .

Let K be an algebraically closed field containing \mathbb{E} ; then K also contains \mathbb{F} , so 1 is a polynomial-linear combination of $\{p_j\}$ over K . Again appealing to Theorem 5.6, there can be no common root of $\{p_j\}$ over K ; in particular, there can be no common root over \mathbb{E} . This completes the proof. \square

We are now ready for the main theorem.

Theorem 5.8. *Let p equal 0 or a prime and let \mathbb{F} be the associated characteristic field. For any graph G , there exists a finite-degree extension K of \mathbb{F} with $\text{dp}_K(G) = \hat{\text{d}}\text{p}_p(G)$.*

Proof. Let \mathbb{E} be a field with characteristic p such that $\text{dp}_{\mathbb{E}}(G) = \hat{\text{d}}\text{p}_p(G)$. Suppose G has n vertices and recall the indicator $\chi_{i,j}$, which is 1 if vertices i and j are adjacent and 0 otherwise.

Let $d = \hat{\text{d}}\text{p}_p(G)$ and define the vectors (whose entries are variables) $v_j = (x_1^j, \dots, x_d^j)$ for each $j = 1, 2, \dots, n$. Then

$$\{v_i \cdot v_j = \chi_{i,j}\}_{1 \leq i < j \leq n}$$

is a system of algebraic equations in the finite variable set $\{x_i^j\}$ which has a solution over \mathbb{E} .

Further, the coefficients of these equations are all 1 or 0, so lie in \mathbb{F} . Applying Proposition 5.7, there exists a finite-degree extension K of \mathbb{F} containing a solution to this system; such a solution gives a d -dimensional representation, so $\text{dp}_K(G) \leq d = \hat{\text{d}}_p(G)$. K has characteristic p , so $\text{dp}_K(G) \geq \hat{\text{d}}_p(G)$, completing the proof. \square

For a prime p , $\hat{\text{d}}_p(G)$ is thus achieved by some finite field \mathbb{F}_{p^e} ; $\hat{\text{d}}_0(G)$ is achieved by some finite-degree extension of \mathbb{Q} . Every such extension is contained in \mathbb{C} , so $\hat{\text{d}}_0(G) = \text{dp}_{\mathbb{C}}(G)$ for every graph G . Recall, though, that $\text{dp}_{\mathbb{C}}$ refers to dot product representations over \mathbb{C} , not Hermitian representations; it is not true that $\hat{\text{d}}_0(G) = \text{dp}^+(G)$ (the right side, recall, equals $\text{dp}_{\mathbb{R}}(G)$) for every graph. For one example, the star graph S_n on n vertices is a tree, so has $\text{dp}_{\mathbb{R}}(S_n) = n - 1$, but $\hat{\text{d}}_0(S_n) = \text{dp}_{\mathbb{C}}(S_n) = 2$ (assign the root the vector $(1, 0)$ and all the leaves $(1, i)$).

Notice that the proof of Theorem 5.8 allows us to reduce the problem of finding the limiting dimension to checking whether a set of polynomials generates a proper ideal over the characteristic field. More specifically, we can check if $\hat{\text{d}}_p(G) \leq d$ for some d by writing the $\binom{n}{2}$ equations $v_i \cdot v_j = \chi_{i,j}$ (on dn variables), and asking whether or not 1 is in the ideal they generate. By generating a Groebner basis, this can be checked by a computer in finite time; thus we actually have an algorithm (albeit a slow one) for computing $\hat{\text{d}}_p(G)$.

Now suppose $\mathbb{Q}[\alpha_1, \dots, \alpha_m]$ is a finite-degree extension over which a graph G achieves its limiting dimension. A given dot product representation of this limiting dimension can thus be written using a finite number of rationals. If we choose a prime p which is not in the denominator of any of these rationals, then this representation passes to a representation of G over some finite field of characteristic p . Thus for any graph G , $\hat{\text{d}}_p(G) \leq \hat{\text{d}}_0(G)$ for all but finitely many primes p .

Another property of limiting dimensions relates to singularity. Singularity implies, in a sense, that there “should” be a representation of smaller dimension, but the field just lacks the elements to express it. We thus might expect for singularity to disappear when the fields become “big enough.” After a preparatory lemma, we give a theorem expressing this.

Lemma 5.9. *Let \mathbb{E} be a finite-degree extension of \mathbb{F} . If A is a matrix over \mathbb{F} with rank r , then its rank is also r when viewed as a matrix over \mathbb{E} .*

Proof. Pick a basis for the column space of A (as a matrix over \mathbb{F}); it is still linearly independent over \mathbb{E} by Proposition 5.4. Since it still spans the

column space, it forms a basis, so the rank is still r . \square

Theorem 5.10. *Let $p > 2$ be prime, G a graph, and \mathbb{F} a finite field of characteristic p such that $\text{dp}_{\mathbb{F}}(G) = \hat{\text{d}}_p(G)$. G is not singular over \mathbb{F} .*

Proof. Suppose G is singular over \mathbb{F} . From Theorem 4.22, there exists a diagonal matrix D with $\text{rank}(A + D) = \text{dp}_{\mathbb{F}}(G) - 1$. Let this rank be r ; by Proposition 4.3, there exists an $n \times r$ matrix T and a symmetric, invertible $r \times r$ matrix U such that $A + D = TUT^T$. Since U is an invertible matrix over \mathbb{F} , its determinant is some nonzero element α of \mathbb{F} . Now $\mathbb{E} = \mathbb{F}[\sqrt{\alpha}]$ is a finite field in which $\det U = \alpha$ is a square. By Lemma 5.9, U is still invertible over \mathbb{E} ; thus the factorization $A + D = TUT^T$ is valid to define χ , so $\chi_{\mathbb{E}}(A + D) = 1$. By Theorem 4.9, $\text{dp}_{\mathbb{E}}(G) \leq \text{rank}(A + D) = r < \text{dp}_{\mathbb{F}}(G)$. Thus $\text{dp}_{\mathbb{F}}(G)$ cannot be $\hat{\text{d}}_p(G)$, completing the proof by contrapositive. \square

This result gives us a family of graphs for which $\hat{\text{d}}_p(G) < \text{dp}_{\mathbb{F}_p}(G)$ (for $p > 2$): namely, any singular graph. In the characteristic two case, singularity is associated with the zero-diagonal adjacency matrix having minimum rank. The condition of having all zeros on the diagonal is not changed by moving to a field extension, so we cannot expect a proof like the one just given to work. In fact, as shown by an example in the following section, the result does not hold.

5.2 Example Graphs

The star graphs $S_n = K_{1,n-1}$ provide examples of singular graphs over fields which do not contain $\sqrt{-1}$ (like \mathbb{F}_p for primes $p \equiv 3 \pmod{4}$). Further, as guaranteed by Theorem 5.10, they are no longer singular in an appropriate field extension. In the following few results we prove the value of the dot product dimension for the star graph over finite fields with characteristic larger than two. (Notice at the outset that the star graph on more than 2 vertices cannot have dot product dimension 1 over any field, since it is not a single clique.)

Lemma 5.11. *Let $p > 2$ be prime. There exist nonzero x, y with $x^2 + y^2 \equiv 0 \pmod{p}$ iff $p \equiv 1 \pmod{4}$. If $p \equiv 3 \pmod{4}$ then there exist nonzero x, y, z with $x^2 + y^2 + z^2 \equiv 0 \pmod{p}$.*

Proof. It is well-known that -1 is a quadratic residue mod p iff $p \equiv 1 \pmod{4}$. If $p \equiv 1$ and $i^2 \equiv -1$, then $(x, y) = (i, 1)$ suffices. Conversely, if $x^2 + y^2 \equiv 0$

with $y \neq 0$, then $(xy^{-1})^2 \equiv -1$ so $p \equiv 1 \pmod{4}$. This proves the first statement.

For the second, recall that there are $(p-1)/2$ nonzero quadratic residues modulo p . Since there do not exist nonzero x, y with $x^2 \equiv -y^2$, the set $\{x^2 : x \neq 0\}$ of $(p-1)/2$ squares is disjoint from the set $\{-y^2 : y \neq 0\}$, which also has $(p-1)/2$ values. Since there are only $p-1$ total nonzero elements, each must be either a square or the negative of a square.

Now suppose for contradiction that all sums of two squares are squares; that is, for all $x, y \neq 0$, there exists z such that $x^2 + y^2 \equiv z^2$. There are $(p-1)/2$ values for the left-hand side for a given x , which implies that the equation ranges over all nonzero quadratic residues for each given x . Thus, there exists $y \neq 0$ such that $x^2 + y^2 \equiv x^2$; this is a contradiction, so there indeed exists some sum of two squares $x^2 + y^2$ which is not a square. Then by the argument of the previous paragraph, $x^2 + y^2 \equiv -z^2$ for some z , so $x^2 + y^2 + z^2 \equiv 0$, as desired. \square

Lemma 5.12. *Let $p > 2$ be prime and pick some $n \geq 4$. $\text{dp}_{\mathbb{F}_p}(S_n)$ equals 2 if $p \equiv 1 \pmod{4}$ and equals 3 otherwise. In either case, $\text{dp}_{\mathbb{F}_{p^2}}(S_n) = 2$.*

Proof. Suppose $p \equiv 1$. Then -1 is a square in \mathbb{F}_p , so there exists i with $i^2 = -1 \implies i^2 + 1^2 = 0$. Now consider the two-dimensional representation where we assign to the root the basis vector $(1, 0)$ and to each other vertex the vector $(1, i)$. Dot products are easily verified to see that this is indeed a representation. Thus $\text{dp}_{\mathbb{F}_p}(S_n) \leq 2$; since the dot product dimension is larger than one, $\text{dp}_{\mathbb{F}_p}(S_n) = 2$.

Now suppose $p \equiv 3$. Then by Lemma 5.11 there do not exist $x, y \neq 0$ with $x^2 + y^2 = 0$ in \mathbb{F}_p . Thus, in particular, there are no nonzero two-dimensional vectors (x, y) which are self-orthogonal. Let v_1, v_2, v_3 be vertices other than the root r , and suppose for contradiction there exists a representation X of dimension two. Then $X(v_1)$ cannot be self-orthogonal and $X(v_1) \cdot X(v_2) = 0$, so $X(v_1)$ and $X(v_2)$ are linearly independent. Further, $X(r)$ must be linearly independent from both of them, as $X(r) \cdot X(v_3) = 1$ but $X(v_1) \cdot X(v_3) = 0$. We cannot have a linearly independent set of size 3 in a space of dimension 2, providing the contradiction. Thus there do not exist representations of dimension 2, so $\text{dp}_{\mathbb{F}_p}(S_n) \geq 3$. Again appealing to Lemma 5.11, there exist nonzero x, y, z with $x^2 + y^2 + z^2 = 0$ in \mathbb{F}_p ; assigning $(x^{-1}, 0, 0)$ to the root and (x, y, z) to all other vertices gives a three-dimensional representation. Together with the above inequality, this shows $\text{dp}_{\mathbb{F}_p}(S_n) = 3$.

Finally, in the case where -1 is not a quadratic residue, $\mathbb{F}_p[\sqrt{-1}]$ is a two-dimensional extension, so by uniqueness must be \mathbb{F}_{p^2} . In this field there do exist nonzero x, y with $x^2 + y^2 = 0$, so $(x^{-1}, 0)$ for the root and (x, y) for all others defines a two-dimensional representation. \square

It is interesting to observe that the star graph behaves very differently over finite fields and the reals; since it is a tree, $\text{dp}_{\mathbb{R}}(S_n) = n - 1$. (Thus $\text{dp}_{\mathbb{Q}}(S_n) = n - 1$ as well.) This stands in stark contrast with the result just proved, that $\text{dp}_{\mathbb{F}}(S_n) \leq 3$ for any finite field \mathbb{F} .

Collecting the results above, we have shown the following proposition.

Proposition 5.13. *Pick $n \geq 4$. Let p be a prime congruent to 3 mod 4. Then S_n is singular over \mathbb{F}_p with dot product dimension 3. For any prime p , the limiting dimension is $\hat{\text{dp}}_p(S_n) = 2$.*

As we claimed above, it is not true that a graph achieving its limiting dimension for characteristic two must be nonsingular. Consider the graph on five vertices $\overline{2K_2}$, shown in Figure 5.1, which is our running example of a singular graph over \mathbb{F}_2 .




Figure 5.1: A singular graph over \mathbb{F}_2

For convenience, name this graph S_* ; a computational proposition shows that S_* is *always* singular.

Proposition 5.14. *S_* is singular over every finite field of characteristic two.*

Proof. Let \mathbb{F} be a finite field of characteristic two. Then \mathbb{F} contains \mathbb{F}_2 , so in particular it contains the deficient representation of dimension 3 which we used to prove S_* is singular over \mathbb{F}_2 . This representation is still deficient over \mathbb{F} as ranks are unchanged by Lemma 5.9. Thus, if we assume for contradiction that S_* is not singular over \mathbb{F} , then we must have $\text{dp}_{\mathbb{F}}(S_*) < 3$ so that this deficient representation does not have minimum dimension. Let A be the adjacency matrix of S_* .

If $\text{dp}_{\mathbb{F}}(S_*) < 3$, then there exists a nonzero diagonal matrix D such that $\text{rank}(A + D) = \text{dp}_{\mathbb{F}}(S_*) < 3$, so $\text{rank}(A + D) \leq 2$. Let $D = \text{diag}(x_1, \dots, x_5)$.

Then

$$A + D = \begin{pmatrix} x_1 & 1 & 1 & 1 & 1 \\ 1 & x_2 & 0 & 1 & 1 \\ 1 & 0 & x_3 & 1 & 1 \\ 1 & 1 & 1 & x_4 & 0 \\ 1 & 1 & 1 & 0 & x_5 \end{pmatrix}.$$

We can transform this matrix by elementary row operations (which do not change the rank) to

$$\begin{pmatrix} x_1 & x_2 & 0 & x_4 & 0 \\ 0 & x_2 & x_3 & 0 & 0 \\ 0 & 0 & 0 & x_4 & x_5 \\ 1 & 0 & x_3 & 1 & 1 \\ 1 & 1 & 1 & x_4 & 0 \end{pmatrix},$$

where we used several times that $1 + 1 = 0$, but did not require any other properties of the field. Notice that the bottom two rows are always linearly independent, as their restrictions to the second and fifth columns are linearly independent. Thus $\text{rank}(A + D) \geq 2$, so we must have $\text{rank}(A + D) = 2$. However, this means that each of the top three rows may be written as a linear combination of the bottom two. The only combination which could give the second row is

$$(0, x_2, x_3, 0, 0) = x_2(1, 0, x_3, 1, 1) + x_2(1, 1, 1, x_4, 0)$$

by looking at the first and second columns. Looking at the last coordinate, $x_2 = 0$ so the right-hand side is zero; thus $x_3 = 0$ as well. Now the only linear combination which could give the third row is

$$(0, 0, 0, x_4, x_5) = x_5(1, 0, x_3, 1, 1)$$

by looking at the second and last columns. Thus $x_5 = 0$, so the right-hand side is zero and thus $x_4 = 0$. Substituting these values, the first row $(x_1, 0, 0, 0, 0)$ must be a linear combination of the bottom two rows $(1, 0, 0, 1, 1)$ and $(1, 1, 1, 0, 0)$; this is only possible if $x_1 = 0$.

We have shown $D = \text{diag}(x_1, \dots, x_5) = 0$, but D was assumed nonzero. This contradicts our opening assumption, and proves that S_* is indeed singular. \square

Singularity provided us with examples of graphs for which $\hat{\text{dp}}_p(G) < \text{dp}_{\mathbb{F}_p}(G)$ for $p > 2$, but it does not (in general) for $p = 2$. Indeed, an exhaustive computer search looking for examples found that $\text{dp}_{\mathbb{F}_2}(G) =$

$\text{dp}_{\mathbb{F}_4}(G) = \text{dp}_{\mathbb{F}_8}(G)$ for all graphs on 5 or fewer vertices. However, the same computer search found that the graph in Figure 5.2 has $\text{dp}_{\mathbb{F}_2}$ equal to 4 but $\text{dp}_{\mathbb{F}_4}$ equal to 3. (This graph has diameter 3, so $\hat{\text{dp}}_2$ in fact equals 3.)




Figure 5.2: A graph G with $\text{dp}_{\mathbb{F}_4}(G) \neq \text{dp}_{\mathbb{F}_2}(G)$

In summary of this discussion, for all prime characteristics p , there are graphs which achieve their limiting dimension over \mathbb{F}_p (the path, for example), and there are also graphs which do not. However, the case $p = 2$ is substantively different in that its limiting dimension *can* be achieved by a deficient representation; singularity over a field achieving the limiting dimension is not possible in any other prime characteristic.

Chapter 6

Conclusions

In this thesis we studied dot product representations over fields other than \mathbb{R} ; in particular, over the complex numbers, rational numbers, and finite fields. We looked at these with the broad mandate of understanding their similarities and differences with the real case. To that end, we proved that the complex case is identical to the real case (if we consider Hermitian representations), and that the rational numbers may differ in exact dimension but always have the same asymptotic dimension as the real numbers. We also gave a matrix characterization, in clear analogue with a preexisting result over \mathbb{R} , of the dot product dimension over finite fields, and studied a few properties of the dot product dimension over finite fields which do not arise in the real case. We gave a characterization of the dot product dimension in the special case \mathbb{F}_2 using the idea of graph toggling, and studied some basic behavior under field extensions.

However, there are a variety of different directions to pursue for future research. We close this report by surveying some of the possibilities.

For one, there is embarrassingly little known about the (exact) dot product dimension over the rational numbers; future research should look further at this case. In particular, we do not even know the dot product dimension for the wheel graph W_6 . In Chapter 3 we mentioned that it is larger than 3, but whether the actual dimension is 4 or 5 is unknown.

Another direction suggested obviously by this work is to look at sub-additivity and leaf removal over arbitrary finite fields. Further, even in the most-studied case of \mathbb{F}_2 we do not have a satisfying characterization of the dot product dimension of trees.

In the finite field cases, we gave exponential-time algorithms for computing the dot product dimension. This leaves open the question of whether

a better (faster) algorithm is possible; it would be interesting to study these problems from a complexity point of view. In particular, perhaps finding the dot product dimension over a finite field is an NP-complete problem.

Our study of limiting dimensions in Chapter 5 was rather brief, and several questions remain in that avenue of research. The only known general bound on the limiting dimension is the diameter of the graph; relationships with other quantities have not been investigated. Also, we made brief mention of the relationship between the limiting dimension of characteristic zero and of prime characteristic. A formal study of this may yield interesting results.

In a broader sense, we have only been considering representations over fields with the standard dot product (or, in the complex case, the Hermitian inner product). This limitation has good reason; fields with the dot product are intimately connected to matrix theory, which allows us to bring many powerful tools to bear on the problem. However, it would be interesting to see which properties of the dot product dimension survive a change to (1) arbitrary bilinear forms or (2) rings (like $\mathbb{Z}/n\mathbb{Z}$) instead of fields.

Overall, the theory of dot product dimensions is a young field which has many possibilities for new research. Hopefully this thesis has helped convince the reader of the motivation for and interest in such work.

Bibliography

- Barrett, Wayne, Jason Grout, and Raphael Loewy. 2006. The minimum rank problem over the finite field of order 2: Minimum rank 3. arXiv:math/0612331.
- Chenette, Nathan L., Sean V. Droms, Leslie Hogben, Rana Mikkelsen, and Olga Pryporova. 2007. Minimum rank of a tree over an arbitrary field. *Electronic Journal of Linear Algebra* 16:183–186.
- Dummit, David S., and Richard M. Foote. 1999. *Abstract Algebra*. John Wiley and Sons, 2nd ed.
- Horn, Roger A., and Charles R. Johnson. 1990. *Matrix Analysis*. Cambridge: Cambridge University Press.
- Laurent, Monique. 2001. Matrix completion problems. *The Encyclopedia of Optimization* 3:221–229.
- Lovász, László, Michael E. Saks, and Alexander Schrijver. 1989. Orthogonal representations and connectivity of graphs. *Linear Algebra Applications* 114/115:439–454.
- Lovász, László, and Katalin Vesztergombi. 1999. Geometric representations of graphs. In Paul Erdos, Proc. Conf., Budapest, 1999. URL <http://microsoft.com/users/lovasz/egrep.ps>.
- MacWilliams, Jessie. 1969. Orthogonal matrices over finite fields. *American Mathematical Monthly* 76(2):152–164.
- Nickel, Christine. 2007. *Random Dot Product Graphs: A Model for Social Networks*. PhD dissertation, Johns Hopkins University, Department of Applied Mathematics and Statistics.

- Peeters, René. 1996. Orthogonal representations over finite fields and the chromatic number of graphs. *Combinatorica* 16(3):417–431.
- Tucker, Kimberly. 2007. *Exact and Asymptotic Dot Product Representations of Graphs*. Ph.D. thesis, Johns Hopkins University, Department of Applied Mathematics and Statistics.
- Tucker, Kimberly, and Edward R. Scheinerman. 2006a. Exact and asymptotic dot product representations of graphs I: Fundamentals. Submitted.
- . 2006b. Exact and asymptotic dot product representations of graphs II: Characterization and recognition. Submitted.