Chapter 6

Orthogonal representations: the smallest cone

In this chapter we discuss the quantity $\vartheta(G)$ we get by asking for the smallest cone in which a given graph G has an orthogonal representation. This quantity will be related to stability number, clique number and chromatic number of graphs.

We describe several applications of this graph parameter, of which we mention two problems in which this parameter yields the only known solution. One of these is the problem of the Shannon capacity (or zero-error capacity) of a graph, which is a tough parameter to compute, and whose evaluation even for a small graph like the pentagon needs the theta function. The other is the computation of clique number and chromatic number for perfect graphs in polynomial time.

6.1 Orthogonal representations and the theta function

In this chapter we study representations that are "economical" in a different sense. It turns out that the smallest half-angle ϕ of a rotational cone (in arbitrary dimension) which contains all vectors in an orthogonal representation of the graph contains interesting information about the graph.

To be precise, we will work with a transformed version of this quantity, namely

$$\vartheta(G) = \frac{1}{(\cos\phi)^2} = \min_{\mathbf{u},\mathbf{c}} \max_{i \in V} \frac{1}{(\mathbf{c}^{\mathsf{T}}\mathbf{u}_i)^2},$$

where the minimum is taken over all orthonormal representations $(\mathbf{u}_i : \ i \in V)$ of G and all unit vectors $\mathbf{c}$. We call $\mathbf{c}$ the "handle" of the representation. Of course, we could fix $\mathbf{c}$ to be (say) the standard basis vector $\mathbf{e}_1$, but this is not always convenient.

The following rather easy inequalities will nevertheless be very important.

Theorem 6.1.1 *For every graph G,*

$$\alpha(G) \le \vartheta(G) \le \chi(\overline{G}).$$

Proof. First, let $S \subseteq V$ be a maximum stable set of nodes in G. Then in every orthonormal representation $(\mathbf{u}_i)$, the vectors $\{\mathbf{u}_i : i \in S\}$ are mutually orthogonal unit vectors. Hence

$$1 = \mathbf{c}^{\mathsf{T}}\mathbf{c} \ge \sum_{i \in S} (\mathbf{c}^{\mathsf{T}}\mathbf{u}_i)^2 \ge |S| \min_i (\mathbf{c}^{\mathsf{T}}\mathbf{u}_i)^2, \tag{6.1}$$

and so

$$\max_{i \in V} \frac{1}{(\mathbf{c}^{\mathsf{T}}\mathbf{u}_i)^2} \ge |S| = \alpha(G).$$

This implies the first inequality. The second follows from the orthogonal representation obtained in Example 5.0.5, using $\mathbf{c} = \frac{1}{\sqrt{m}}(\mathbf{e}_1 + \cdots + \mathbf{e}_m)$ as the handle. □

Our examples of orthogonal representations give further upper bounds for ϑ. From the trivial orthogonal representation we get the trivial bound $\vartheta(G) \le n$. Example 5.0.3 leads to the following construction: Assuming that there are no isolated nodes, we assign to each node $i \in V$ the vector $\mathbf{u}_i = (1/\sqrt{d_i})\nabla_i \in \mathbb{R}^E$. As a handle, take the vector $\mathbf{c} = (1/\sqrt{m})\mathbb{1} \in \mathbb{R}^E$. Then $\mathbf{c}^{\mathsf{T}}\mathbf{u}_i = \sqrt{d_i/m}$, and so we get the bound

$$\vartheta(G) \le \max_i \frac{m}{d_i} = \frac{m}{d_{\min}}.$$

The upper bound $m/d_{\min}$ for the independence number $\alpha(G)$ is easy to prove by counting edges. From Example 5.0.6 we get by elementary trigonometry that

$$\vartheta(C_5) \le \sqrt{5}. \tag{6.2}$$

Soon we'll see that equality holds here.

It is clear that if G' is an induced subgraph of G, then $\vartheta(G') \le \vartheta(G)$ (using that an optimal orthonormal representation of G, restricted to $V(G')$, is an orthonormal representation of G'). It is also clear that if G' is a spanning subgraph of G (i.e., $V(G') = V(G)$ and $E(G') \subseteq E(G)$), then $\vartheta(G') \ge \vartheta(G)$ (using that an optimal orthonormal representation of G' is an orthonormal representation of G).

6.1.1 Duality for theta

The graph parameter ϑ has many equivalent definitions. We are going to state some, which lead to an important dual formulation of this quantity.

Vector chromatic number. The following geometric definition was proposed by Karger, Motwani and Sudan. In terms of the complementary graph, this value is called the "vector

chromatic number". As a motivation for this name, consider a t-colorable graph ($t \geq 2$), and let us color its nodes by t unit vectors $\mathbf{f}_1, \ldots, \mathbf{f}_t \in \mathbb{R}^{t-1}$, pointing to the vertices of a regular simplex. We get a vector labeling $\mathbf{w}$ that is not an orthogonal representation of $\overline{G}$, but it is closely related. It is not hard to compute that $\mathbf{f}_k^{\mathsf{T}}\mathbf{f}_l = -1/(t-1)$ for $k \neq l$, and so $\mathbf{w}_i^{\mathsf{T}}\mathbf{w}_j = -1/(t-1)$ for $ij \in E$.

Now let us forget about the condition that $\mathbf{w}_i$ must be one of the vectors $\mathbf{f}_k$: define a (strict) *vector t-coloring* ($t > 1$) of the graph G as a vector labeling $i \mapsto \mathbf{w}_i \in \mathbb{R}^n$ such that $|\mathbf{w}_i| = 1$ for all $i \in V$, and $\mathbf{w}_i^{\mathsf{T}}\mathbf{w}_j = -1/(t-1)$ for all $ij \in E$. The smallest $t \geq 1$ for which the graph G has a vector t-coloring is called its (strict) *vector chromatic number*, and is denoted by Vchr.

The dimension n in the definition above is just chosen to be large enough; allowing a higher dimension would not make any difference. The definition of vector t-coloring is meaningful for every real number $t > 1$. We may even consider it meaningful for $t = 1$, when it means that $E = \emptyset$, so G is edgeless and $\chi(G) = 1$.

If G is not edgeless, then, trivially, a vector t-coloring of G can exist for $t \geq 2$ only. For $t = 2$, a vector 2-coloring is necessarily strict, and the endpoints of any edge must be labeled by antipodal unit vectors. It follows that for every connected component of G is labeled by two antipodal vectors only, and G is bipartite.

It is clear from the construction above that $\mathsf{Vchr}(G) \leq \chi(G)$ for every graph G. Equality does not hold in general. Labeling the nodes of a pentagon by the vertices of a regular pentagon inscribed in the unit circle, so that the edges are mapped onto the diagonals, we see that $\mathsf{Vchr}(C_5) \leq \sqrt{5} < \chi(C_5) = 3$. It is also easy to see that $\mathsf{Vchr}(G) \geq \omega(G)$.

Semidefinite optimization. Next, we give a couple of formulas for ϑ in terms of semidefinite matrices. Let

$$\vartheta_{\mathrm{mindiag}} = \min\bigl\{1 + \max_{i \in V} Y_{ii} : \; Y \in \mathbb{R}^{V \times V},\; Y \succeq 0,\; Y_{ij} = -1 \;(ij \in \overline{E})\bigr\} \tag{6.3}$$

and

$$\vartheta_{\mathrm{maxsum}} = \max\Bigl\{\sum_{i,j \in V} Z_{ij} : \; Z \in \mathbb{R}^{V \times V},\; Z \succeq 0,\; Z_{ij} = 0 \;(ij \in E),\; \mathrm{tr}(Z) = 1\Bigr\}. \tag{6.4}$$

It will turn out that these two values are equal. This equality is in fact a special case of the Duality Theorem of semidefinite optimization. It is not hard to check that (6.3) and (6.4) are dual semidefinite programs, and the first one has a strictly feasible solution. So the Duality Theorem of semidefinite programming applies, and asserts that the two programs have the same objective value. However, we are going to include a proof, to make our treatment self-contained.

Dual orthogonal representation. We use orthonormal representations of the complemen-

tary graph to define

$$\vartheta_{\rm dual} = \max \sum_{i \in V} (\mathbf{d}^{\mathsf{T}} \mathbf{v}_i)^2, \tag{6.5}$$

where the maximum extends over all orthonormal representations $(\mathbf{v}_i : i \in V)$ of the complementary graph $\overline{G}$ and all unit vectors (handles) $\mathbf{d}$.

The main theorem of this section asserts that all these definitions lead to the same value.

Theorem 6.1.2 *For every graph G, $\vartheta(G) = \mathsf{Vchr}(\overline{G}) = \vartheta_{\rm mindiag}(G) = \vartheta_{\rm maxsum}(G) = \vartheta_{\rm dual}(G)$.*

Proof. We prove the circle of inequalities

$$\vartheta(G) \le \mathsf{Vchr}(\overline{G}) \le \vartheta_{\rm mindiag}(G) \le \vartheta_{\rm maxsum}(G) \le \vartheta_{\rm dual}(G) \le \vartheta(G). \tag{6.6}$$

To prove the first inequality, let $t = \mathsf{Vchr}(\overline{G})$, and let $(\mathbf{w}_i : i \in V)$ be an optimal vector t-coloring. Let $\mathbf{c}$ be a vector orthogonal to all the $\mathbf{w}_i$ (we increase the dimension of the space if necessary). Let

$$\mathbf{u}_i = \frac{1}{\sqrt{t}}\mathbf{c} + \sqrt{\frac{t-1}{t}}\mathbf{w}_i.$$

Then $|\mathbf{u}_i| = 1$ and $\mathbf{u}_i^{\mathsf{T}}\mathbf{u}_j = 0$ for $ij \in \overline{E}$, so $(\mathbf{u}_i)$ is an orthonormal representation of G. Furthermore, with handle $\mathbf{c}$ we have $\mathbf{c}^{\mathsf{T}}\mathbf{u}_i = 1/\sqrt{t}$, which implies that $\vartheta(G) \le \mathsf{Vchr}(\overline{G})$.

Second, let Y be an optimal solution of (6.3). We may assume that all diagonal entries Y_{ii} are the same number t, since we can replace all of them by the largest without violating the other constraints. The matrix $1/(t-1)Y$ is positive semidefinite, and so it can be written as $\mathsf{Gram}(\mathbf{w}_i : i \in V)$ with appropriate vectors $(\mathbf{w}_i \in \mathbb{R}^n)$. These vectors form a strict vector t-coloring. Since $\mathsf{Vchr}(\overline{G})$ is the smallest t for which this exists, this proves that $\mathsf{Vchr}(\overline{G}) \le t = \vartheta_{\rm mindiag}$.

The main step in the proof is to show that $\vartheta_{\rm mindiag} \le \vartheta_{\rm maxsum}$. Fix any $t > \vartheta_{\rm maxsum}$; it is easy to see that $\vartheta_{\rm maxsum} \ge 1$ and hence $t > 1$. Let $\mathcal{L}_t$ denote the linear space of symmetric $V \times V$ matrices satisfying $Z_{ij} = 0$ $(ij \in E)$ and $(tI - J) \cdot Z = 0$, and let $\mathcal{P}_V$ denote the cone of positive semidefinite $V \times V$ matrices.

We claim that $\mathcal{P}_V \cap \mathcal{L}_t = \{0\}$. Suppose, to the contrary, that there is a symmetric matrix $Z \neq 0$ such that $Z \in \mathcal{P}_V \cap \mathcal{L}_t$. Every nonzero positive semidefinite matrix has positive trace, so $\mathrm{tr}(Z) > 0$; by scaling, we may assume that $\mathrm{tr}(Z) = 1$. Then Z satisfies the conditions in the definition of $\vartheta_{\rm maxsum}$, and so $\vartheta_{\rm maxsum} \ge J \cdot Z = tI \cdot Z = t$, contradicting the choice of t.

It follows that there is a hyperplane $\mathcal{H}$ through the origin such that $\mathcal{L}_t \subseteq \mathcal{H}$ and $\mathcal{H} \cap \mathcal{P}_V = \{0\}$. Let $Y \cdot X = 0$ be the equation of $\mathcal{H}$ (where $Y \neq 0$ is a symmetric $V \times V$ matrix); we may assume that $Y \cdot X \ge 0$ for all $X \in \mathcal{P}_V$. This means that Y is in the polar cone of $\mathcal{P}_V$,

which means that $Y \succeq 0$. Furthermore, $\mathcal{L}_t \subseteq \mathcal{H}$ means that the equation of $\mathcal{H}$ is a linear combination of the equations defining $\mathcal{L}_t$, i.e., there are real numbers a_{ij} ($ij \in E$) and b such that

$$Y = \sum_{ij \in E} a_{ij} E_{ij} + b(tI - J).$$

Considering a positive diagonal entry of Y, we see that $b > 0$, and since we are free to scale Y by positive scalars, we may assume that $b = 1$. But this means that Y satisfies the conditions in the definition of $\vartheta_{\mathrm{mindiag}}$, and so $\vartheta_{\mathrm{mindiag}} \le 1 + \max_i Y_{ii} = t$. Since this holds for every $t > \vartheta_{\mathrm{maxsum}}$, this implies that $\vartheta_{\mathrm{mindiag}} \le \vartheta_{\mathrm{maxsum}}$.

To prove the fourth inequality in (6.6), let Z be an optimum solution of the program (6.4) with objective function value $\vartheta_{\mathrm{maxsum}}$. We can write Z as $\mathsf{Gram}(\mathbf{z}_i : i \in V)$ where $\mathbf{z}_i \in \mathbb{R}^n$. Let us rescale the vectors $\mathbf{z}_i$ to get the unit vectors $\mathbf{v}_i = \mathbf{z}_i^0$ (if $\mathbf{z}_i = 0$ then we take a unit vector orthogonal to everything else as $\mathbf{v}_i$). Define $\mathbf{d} = (\sum_i \mathbf{z}_i)^0$.

By the properties of Z, the vectors $\mathbf{v}_i$ form an orthonormal representation of $\overline{G}$, and hence

$$\vartheta_{\mathrm{dual}} \ge \sum_i (\mathbf{d}^{\mathsf{T}} \mathbf{v}_i)^2.$$

To estimate the right side, we use the equations

$$\sum_i |\mathbf{z}_i|^2 = \sum_i \mathbf{z}_i^{\mathsf{T}} \mathbf{z}_i = \mathrm{tr}(Z) = 1, \qquad \Big|\sum_i \mathbf{z}_i\Big|^2 = \sum_{i,j} \mathbf{z}_i^{\mathsf{T}} \mathbf{z}_j = \sum_{i,j} Z_{ij} = \vartheta_{\mathrm{maxsum}},$$

and the Cauchy–Schwarz Inequality:

$$\begin{aligned}\sum_i (\mathbf{d}^{\mathsf{T}} \mathbf{v}_i)^2 &= \Big(\sum_i |\mathbf{z}_i|^2\Big)\Big(\sum_i (\mathbf{d}^{\mathsf{T}} \mathbf{v}_i)^2\Big) \ge \Big(\sum_i |\mathbf{z}_i| \mathbf{d}^{\mathsf{T}} \mathbf{v}_i\Big)^2 \\ &= \Big(\sum_i \mathbf{d}^{\mathsf{T}} \mathbf{z}_i\Big)^2 = \Big(\mathbf{d}^{\mathsf{T}} \sum_i \mathbf{z}_i\Big)^2 = \Big|\sum_i \mathbf{z}_i\Big|^2 = \vartheta_{\mathrm{maxsum}}.\end{aligned}$$

This proves that $\vartheta_{\mathrm{dual}} \ge \vartheta_{\mathrm{maxsum}}$.

Finally, to prove the last inequality in (6.6), it suffices to prove that if $(\mathbf{u}_i : i \in V)$ is an orthonormal representation of G in $\mathbb{R}^n$ with handle $\mathbf{c}$, and $(\mathbf{v}_i : i \in V)$ is an orthonormal representation of $\overline{G}$ in $\mathbb{R}^m$ with handle $\mathbf{d}$, then

$$\sum_{i \in V} (\mathbf{d}^{\mathsf{T}} \mathbf{v}_i)^2 \le \max_{i \in V} \frac{1}{(\mathbf{c}^{\mathsf{T}} \mathbf{u}_i)^2}. \tag{6.7}$$

The tensor product vectors $\mathbf{u}_i \circ \mathbf{v}_i$ ($i \in V$) are mutually orthogonal unit vectors. Indeed, $(\mathbf{u}_i \circ \mathbf{v}_i)^{\mathsf{T}}(\mathbf{u}_j \circ \mathbf{v}_j) = (\mathbf{u}_i^{\mathsf{T}} \mathbf{u}_j)(\mathbf{v}_i^{\mathsf{T}} \mathbf{v}_j) = 0$, since either $\mathbf{u}_i$ is orthogonal to $\mathbf{u}_j$ or $\mathbf{v}_i$ is orthogonal to $\mathbf{v}_j$. Hence

$$\sum_i (\mathbf{c}^{\mathsf{T}} \mathbf{u}_i)^2 (\mathbf{d}^{\mathsf{T}} \mathbf{v}_i)^2 = \sum_i \big((\mathbf{c} \circ \mathbf{d})^{\mathsf{T}} (\mathbf{u}_i \circ \mathbf{v}_i)\big)^2 \le 1. \tag{6.8}$$

On the other hand,

$$\sum_i (\mathbf{c}^\mathsf{T}\mathbf{u}_i)^2(\mathbf{d}^\mathsf{T}\mathbf{v}_i)^2 \geq \min_i (\mathbf{c}^\mathsf{T}\mathbf{u}_i)^2 \sum_i (\mathbf{d}^\mathsf{T}\mathbf{v}_i)^2,$$

which implies that

$$\sum_i (\mathbf{d}^\mathsf{T}\mathbf{v}_j)^2 \geq \frac{1}{\min_i (\mathbf{c}^\mathsf{T}\mathbf{u}_i)^2} = \max_i \frac{1}{(\mathbf{c}^\mathsf{T}\mathbf{u}_i)^2}.$$

This proves (6.7) and completes the proof of Theorem 6.1.2. □

Remark 6.1.3 We can state the theorem more explicitly as the following sequence of formulas.

$$\vartheta(G) = \min\Big\{\max_{i\in V} \frac{1}{(\mathbf{c}^\mathsf{T}\mathbf{u}_i)^2} : \ \mathbf{u} \text{ ONR of } G,\ |\mathbf{c}| = 1\Big\} \tag{6.9}$$

$$= \min\Big\{t \geq 2 : |\mathbf{w}_i| = 1,\ \mathbf{w}_i^\mathsf{T}\mathbf{w}_j = -\frac{1}{t-1}\ (ij \in \overline{E})\Big\} \tag{6.10}$$

$$= \min\big\{1 + \max_{i\in V} Y_{ii} : \ Y \succeq 0,\ Y_{ij} = -1\ (ij \in \overline{E})\big\} \tag{6.11}$$

$$= \max\Big\{\sum_{i,j\in V} Z_{ij} : \ Z \succeq 0,\ Z_{ij} = 0\ (ij \in E),\ \mathrm{tr}(Z) = 1\Big\} \tag{6.12}$$

$$= \max\Big\{\sum_{i\in V} (\mathbf{d}^\mathsf{T}\mathbf{v}_i)^2 : \ \mathbf{v} \text{ ONR of } \overline{G},\ |\mathbf{d}| = 1\Big\}. \tag{6.13}$$

From this form it is clear (and we have seen this in the proof as well) that the powerful step in this sequence of formulas is the equality (6.11)=(6.12), where an expression as a minimum switches to an expression as a maximum. Note that before this equality we have conditions on the edges of $\overline{G}$, which then get replaced by conditions on the edges of G.

6.1.2 Consequences of duality

From the fact that equality holds in (6.6), it follows that equality holds in all of the arguments above. Let us formulate some consequences. Considering the optimal orthogonal representation constructed in the first step of the proof, we get that in squeezing an orthogonal representation into the narrowest possible cone, we may assume that all vectors are on the boundary: Every graph G has an orthonormal representation $(\mathbf{u}_i)$ with handle $\mathbf{c}$ such that for every node i,

$$\mathbf{c}^\mathsf{T}\mathbf{u}_i = \frac{1}{\sqrt{\vartheta(G)}}. \tag{6.14}$$

From the fact that equality must hold in (6.8), and from the derivation of this inequality, we see that for an optimal orthonormal representation $(\mathbf{u}, \mathbf{c})$ and an optimal dual orthonormal

representation $(\mathbf{v}, \mathbf{d})$, the unit vector $\mathbf{c} \circ \mathbf{d}$ must be a linear combination of the mutually orthogonal unit vectors $\mathbf{u}_i \circ \mathbf{v}_i$. The coefficients are easy to figure out, and we get

$$\mathbf{c} \circ \mathbf{d} = \sum_i (\mathbf{c}^\mathsf{T}\mathbf{u}_i)(\mathbf{d}^\mathsf{T}\mathbf{v}_i)(\mathbf{u}_i \circ \mathbf{v}_i), \tag{6.15}$$

or in a matrix form

$$\mathbf{c}\mathbf{d}^\mathsf{T} = \sum_i (\mathbf{c}^\mathsf{T}\mathbf{u}_i)(\mathbf{d}^\mathsf{T}\mathbf{v}_i)\mathbf{u}_i\mathbf{v}_i^\mathsf{T}. \tag{6.16}$$

Applying these matrices to $\mathbf{c}$, we get

$$\mathbf{d} = \sum_i (\mathbf{c}^\mathsf{T}\mathbf{u}_i)^2(\mathbf{d}^\mathsf{T}\mathbf{v}_i)\mathbf{v}_i. \tag{6.17}$$

Using here an optimal orthonormal representation that satisfies (6.14), we get information about an optimal orthogonal representation in the dual definition of $\vartheta(G)$:

$$\sum_{i \in V} (\mathbf{d}^\mathsf{T}\mathbf{v}_i)\mathbf{v}_i = \vartheta(G)\mathbf{d}. \tag{6.18}$$

Lemma 6.1.4 *For every graph G, we have $\vartheta(G)\vartheta(\overline{G}) \geq n$.*

Proof. Let $(\mathbf{u}, \mathbf{c})$ be an optimal orthogonal representation of G. Then applying (6.13) to the complementary graph, we get

$$\vartheta(\overline{G}) \geq \sum_i (\mathbf{c}^\mathsf{T}\mathbf{u}_i)^2 \geq n \min_i (\mathbf{c}^\mathsf{T}\mathbf{u}_i)^2 = \frac{n}{\vartheta(G)}. \qquad \square$$

Equality does not hold in 6.1.4 in general, but it does when G has a node-transitive automorphism group. We say that an orthonormal representation $(\mathbf{u}_i, \mathbf{c})$ in $\mathbb{R}^d$ of a graph G is *automorphism invariant*, if every automorphism $\gamma \in \mathrm{Aut}(G)$ can be lifted to an orthogonal transformation O_γ of $\mathbb{R}^d$ such that $O_\gamma \mathbf{c} = \mathbf{c}$ and $\mathbf{u}_{\gamma(i)} = O_\gamma \mathbf{u}_i$ for every node i. An optimal orthonormal representation (say, in the sense of (6.9)) is not necessarily invariant under automorphisms, but there is always one that is (see Figure 6.1).

Theorem 6.1.5 *Every graph G has an optimal orthonormal representation and an optimal dual orthonormal representation that are both automorphism invariant.*

Proof. We give the proof for the dual orthonormal representation. The optimum solutions of the semidefinite program in (6.3) form a bounded convex set, which is invariant under the transformations $Z \mapsto P_\alpha^\mathsf{T} Z P_\alpha$, where P_α is the permutation matrix defined by an automorphism α of G. If Z is an optimizer in (6.12), then so is $P_\alpha^\mathsf{T} Z P_\alpha$ for every automorphism $\alpha \in \mathrm{Aut}(G)$, and hence also

$$\widehat{Z} = \frac{1}{|\mathrm{Aut}(G)|} \sum_{\alpha \in \mathrm{Aut}(G)} P_\alpha^\mathsf{T} Z P_\alpha.$$

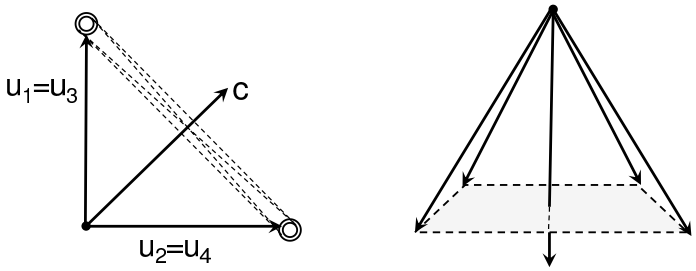

Figure 6.1: An optimal orthonormal representation of C_4 that is not invariant under its automorphisms, and one that is. The representation on the left is also optimal with respect to minimizing the dimension.

This matrix satisfies $P_\alpha^{\mathsf{T}} Z P_\alpha = Z$ for all automorphisms α.

The construction of an orthonormal representation of $\overline{G}$ in the proof of $\vartheta_{\text{mindiag}} \leq \vartheta_{\text{maxsum}}$ in Theorem 6.1.2 can be done in a canonical way: we choose the columns of $Z^{1/2}$ as the vectors $\mathbf{z}_i$, and use them to construct the dual orthonormal representation with $\mathbf{v}_i = \mathbf{z}_i^0$ and $\mathbf{d} = (\sum_i \mathbf{z}_i)^0$. The optimal dual orthonormal representation constructed this way will be invariant under the automorphism group of G. □

Corollary 6.1.6 *If G has a node-transitive automorphism group, then*

$$\vartheta(G)\vartheta(\overline{G}) = n.$$

Proof. It follows from Theorem 6.1.5 that $\overline{G}$ has an orthonormal representation $(\mathbf{v}_i, \mathbf{d})$ in $\mathbb{R}^n$ such that $\sum_i (\mathbf{d}^{\mathsf{T}}\mathbf{v}_i)^2 = \vartheta(G)$, and $\mathbf{d}^{\mathsf{T}}\mathbf{v}_i$ is the same for every i. So $(\mathbf{d}^{\mathsf{T}}\mathbf{v}_i)^2 = \vartheta(G)/n$ for all nodes i, and hence

$$\vartheta(\overline{G}) \leq \max_i \frac{1}{(\mathbf{d}^{\mathsf{T}}\mathbf{v}_i)^2} = \frac{n}{\vartheta(G)}.$$

Since we already know the reverse inequality (Lemma 6.1.4), this proves the Corollary. □

Corollary 6.1.7 *If G is a self-complementary graph with a node-transitive automorphism group, then $\vartheta(G) = \sqrt{n}$. In particular, $\vartheta(C_5) = \sqrt{5}$.*

A further important feature of the theta-function is its nice behavior with respect to graph product; we will see that this is what underlies its applications in information theory.

There are many different ways of multiplying two simple graphs G and H, of which we need one in this chapter. The *strong product* $G \boxtimes H$ is defined on the underlying set $V(G) \times V(H)$. Two nodes (u_1, v_1) and (u_2, v_2) are adjacent if and only if either $ij \in E(G)$ and $uv \in E(H)$, or $ij \in E(G)$ and $u = v$, or $i = j$ and $uv \in E(H)$. It is easy to see that this multiplication is associative and commutative (up to isomorphism). The product of two complete graphs is a complete graph.

Theorem 6.1.8 *For any two graphs G and H, we have $\vartheta(G \boxtimes H) = \vartheta(G)\vartheta(H)$.*

Proof. Let $(\mathbf{u}_i : i \in V)$ be an optimal orthogonal representation of G with handle $\mathbf{c}$ ($\mathbf{u}_i, \mathbf{c} \in \mathbb{R}^n$), and let $(\mathbf{v}_j : j \in V(H))$ be an optimal orthogonal representation of H with handle $\mathbf{d}$ ($\mathbf{v}_j, \mathbf{d} \in \mathbb{R}^m$). It is easy to check that the vectors $\mathbf{u}_i \circ \mathbf{v}_j$ ($(i,j) \in V(G) \times V(H)$) form an orthogonal representation of $G \boxtimes H$. Furthermore, taking $\mathbf{c} \circ \mathbf{d}$ as its handle, we have

$$\big((\mathbf{c} \circ \mathbf{d})^{\mathsf{T}}(\mathbf{u}_i \circ \mathbf{v}_j)\big)^2 = (\mathbf{c}^{\mathsf{T}}\mathbf{u}_i)^2(\mathbf{d} \circ \mathbf{v}_j)^2 \geq \frac{1}{\vartheta(G)} \cdot \frac{1}{\vartheta(H)},$$

and hence

$$\vartheta(G \boxtimes H) \leq \max_{i,j} \frac{1}{\big((\mathbf{c} \circ \mathbf{d})^{\mathsf{T}}(\mathbf{u}_i \circ \mathbf{v}_j)\big)^2} \leq \vartheta(G)\vartheta(H).$$

To prove that equality holds, we use the duality established in Section 6.1.1. Let $(\mathbf{v}_i, \mathbf{d})$ be an orthonormal representation of $\overline{G}$ which is optimal in the sense that $\sum_i (\mathbf{d}^{\mathsf{T}}\mathbf{v}_i)^2 = \vartheta(G)$, and let $(\mathbf{w}_j, \mathbf{e})$ be an orthonormal representation of $\overline{H}$ such that $\sum_i (\mathbf{e}^{\mathsf{T}}\mathbf{w}_i)^2 = \vartheta(H)$. It is easy to check that the vectors $\mathbf{v}_i \circ \mathbf{w}_j$ form an orthonormal representation of $\overline{G \boxtimes H}$, and so using handle $\mathbf{d} \circ \mathbf{e}$ we get

$$\vartheta(G \boxtimes H) \geq \sum_{i,j} \big((\mathbf{d} \circ \mathbf{e})^{\mathsf{T}}(\mathbf{v}_i \circ \mathbf{w}_j)\big)^2 = \sum_{i,j} (\mathbf{d}^{\mathsf{T}}\mathbf{v}_i)^2(\mathbf{e}^{\mathsf{T}}\mathbf{w}_j)^2 = \vartheta(G)\vartheta(H).$$

We already know the reverse inequality, which completes the proof. □

6.1.3 Eigenvalues and theta

To motivate the identities and inequalities to be proved in this section, let us survey some of the classical results which use spectral properties of graphs, or more generally linear algebra techniques, to bound quantities like the independence number $\alpha = \alpha(G)$, the clique number $\omega = \omega(G)$, or the chromatic number $\chi = \chi(G)$. It turns out that several of these results could be used to define $\vartheta(G)$, if generalized appropriately.

Let us start with an almost trivial inequality:

$$\omega \leq 1 + \lambda_{\max}(A_G). \tag{6.19}$$

In terms of the complementary graph,

$$\alpha \leq 1 + \lambda_{\max}(A_{\overline{G}}). \tag{6.20}$$

Indeed, the matrix $A_G + I$ contains an $\omega \times \omega$ submatrix J_ω of 1's, so $1 + \lambda_{\max}(A_G) = \lambda_{\max}(A_G + I) \geq \lambda_{\max}(J_\omega) = \omega$. Note that in this argument, only the matrix entries in adjacent and diagonal positions are used. We could substitute arbitrary real numbers for the remaining entries (which are originally zeroes), to minimize the bound $\lambda_{\max}(A) + 1$. (We

keep A symmetric, to have real eigenvalues.) What is the best bound on ω we can obtain this way?

The following lower bound on the chromatic number of a graph is more difficult to prove:

$$\chi(G) \geq 1 - \frac{\lambda_{\max}(A_G)}{\lambda_{\min}(A_G)} \tag{6.21}$$

(note that $\lambda_{\min}(A_G) < 0$ if G has at least one edge, which we may assume). We will not go through the proof; but if you do, you realize that it uses only the 0's in the adjacency matrix, so we can play with the 1's to get the sharpest possible lower bound. What is the best bound on χ we can obtain this way?

Hoffman (unpublished) proved the following upper bound on α, somewhat analogous to the bound (6.21): If G is a d-regular graph, then

$$\alpha(G) \leq \frac{-n\lambda_{\min}(A_G)}{d - \lambda_{\min}(A_G)} = \frac{-n\lambda_{\min}(A_G)}{\lambda_{\max}(A_G) - \lambda_{\min}(A_G)}. \tag{6.22}$$

Looking at the proof, one realizes that we use here where the 0's of A_G are, and also the fact that all row sums are the same; but not the actual values of the entries corresponding to edges. What is the best bound we can obtain by playing with the entries in adjacent positions?

Perhaps it is not surprising that the answer the first two questions posed above is $\vartheta(\overline{G})$. This will follow from the next identities.

Proposition 6.1.9 *For every graph G,*

$$\vartheta(G) = \min_A \lambda_{\max}(A), \tag{6.23}$$

where A ranges over all $V \times V$-matrices with $A_{ij} = 1$ for $ij \in \overline{E}$ and also for $i = j$. Furthermore,

$$\vartheta(G) = \max_B \lambda_{\max}(B), \tag{6.24}$$

where B ranges over all positive semidefinite $V \times V$-matrices with $B_{ij} = 0$ for $ij \in E$ and $B_{ii} = 1$ for $i \in V$. Furthermore,

$$\vartheta(G) = 1 + \max_C \frac{\lambda_{\max}(C)}{-\lambda_{\min}(C)}, \tag{6.25}$$

where C ranges over all symmetric nonzero $V \times V$-matrices with $C_{ij} = 0$ for $ij \in E$ and also for $i = j$.

Note that (6.24) can be written as

$$\vartheta(G) = \max_{\mathbf{v}} \lambda_{\max}(\mathsf{Gram}(\mathbf{v})), \tag{6.26}$$

where $\mathbf{v}$ ranges over dual orthonormal representations of G.

Proof. By (6.11),

$$\vartheta(G) = \min\bigl\{1 + \max_{i\in V} Y_{ii} : \ Y \succeq 0, \ Y_{ij} = -1 \ (ij \in \overline{E})\bigr\}.$$

Let Y be a minimizer matrix in this formula, let D be the diagonal matrix obtained from Y by changing all off-diagonal entries to 0, and define $A = I + D - Y$. Then A satisfies the conditions in the Proposition, and (using that $Y \succeq 0$)

$$\lambda_{\max}(A) \le \lambda_{\max}(I + D) = 1 + \max_{i\in V} Y_{ii} = \vartheta(G).$$

The reverse inequality follows similarly, by starting with a minimizer in (6.23), and considering $Y = \lambda_{\max}(A)I - A$.

To prove (6.24), we use (6.12):

$$\vartheta(G) = \max\Bigl\{\sum_{i,j\in V} Z_{ij} : \ Z \succeq 0, \ Z_{ij} = 0 \ (ij \in E), \ \mathrm{tr}(Z) = 1\Bigr\}.$$

Let Z be a minimizer here, and define a $V \times V$ matrix B by

$$B_{ij} = \frac{1}{\sqrt{Z_{ii}Z_{jj}}} Z_{ij}.$$

Then B satisfies the conditions in (6.24). Define a vector $\mathbf{x} \in \mathbb{R}^V$ by $x_i = \sqrt{Z_{ii}}$, then $\mathbf{x}$ is a unit vector and

$$\lambda_{\max}(B) \ge \mathbf{x}^{\mathsf{T}} B \mathbf{x} = \sum_{i,j} Z_{ij} = \vartheta(G).$$

The reverse inequality follows similarly, by starting with an optimizer B in (6.24), and scaling its rows and columns by the entries of an eigenvector belonging to $\lambda_{\max}(B)$.

Finally, to prove (6.25), consider an optimizer C in it, then $B = I - \frac{1}{\lambda_{\min}(C)} C$ is positive semidefinite, has $0's$ in adjacent positions and 1's on the diagonal. Hence by (6.24), we have

$$\vartheta(G) \ge \lambda_{\max}(B) = 1 - \frac{\lambda_{\max}(C)}{\lambda_{\min}(C)}.$$

The reverse inequality follows by a similar argument. □

Proposition 6.1.9 can be combined with different known estimates for the largest eigenvalue of a matrix. As an example, using that the largest eigenvalue of a matrix is bounded above by the largest ℓ_1-norm of rows, we get that for every optimal dual orthogonal representation $(\mathbf{v}_i)$,

$$\vartheta(G) \le \max_i \sum_j |\mathbf{v}_i^{\mathsf{T}} \mathbf{v}_j|. \tag{6.27}$$

Finally, we show how a strengthening of Hoffman's bound (6.22) can be derived from Proposition 6.1.9. This gives an upper bound on ϑ in terms of the eigenvalues of the adjacency matrix.

Lemma 6.1.10 *Let G be a d-regular graph. Then for every symmetric nonzero $V \times V$-matrix M such that $M_{ij} = 0$ for $ij \in \overline{E}$ and also for $i = j$, and M has equal row-sums, we have*

$$\vartheta(G) \leq \frac{-n\lambda_{\min}(M)}{\lambda_{\max}(M) - \lambda_{\min}(M)}.$$

If the automorphism group of G is transitive on the nodes, then there is such a matrix M attaining equality. If the automorphism group is transitive on the edges, then equality holds for $M = A_G$.

Proof. Let M have eigenvalues $\lambda_1 = d \geq \lambda_2 \geq \cdots \geq \lambda_n$. The matrix $J - tM$ satisfies the conditions in (6.23) for every value of t. Using the condition that all row-sums of M are the same, we see that $\mathbb{1}$ is a common eigenvector of J and M, and it follows that all eigenvectors of M are eigenvectors of J as well. Hence the eigenvalues of $J - tM$ are $n - td, -t\lambda_2, \ldots, -t\lambda_n$. The largest one is either $n - td$ or $-t\lambda_{\min}$, and we get the best bound if we choose t so that these two are equal: $t = n/(d - \lambda_{\min})$, giving the bound in the lemma.

We can see just as in the proof of Theorem 6.1.5 that there is an optimizing matrix A in (6.23) that is invariant under the automorphisms. So if G has a node-transitive automorphism group, then the row-sums of this matrix are equal, and the same holds for $M = J - A$. This matrix M satisfies the conditions in the lemma, and attains equality.

If G has an edge-transitive automorphism group, then all nonzero entries of M are the same, and hence $M = tA$ for some $t \neq 0$. The value of t cancels from the formula, so $M = A_G$ also provides equality. □

6.1.4 Examples

We compute the theta-function of several classes of graphs, to illustrate the use of the results presented above.

Example 6.1.11 (Cycles) Even cycles are trivial: If n is even, then $\alpha(C_n) = \vartheta(C_n) = \chi(\overline{C}_n) = n/2$. To derive the theta function on odd cycles, we can use Lemma 6.1.10: The eigenvalues of C_n are $2\cos(2k\pi/n)$ $(k = 0, 1, \ldots, n-1)$, of which $k = 0$ gives the largest (which is 2) and $k = (n-1)/2$ gives the smallest one (which is $2\cos((n-1)\pi/n) = -2\cos(\pi/n)$. Hence

$$\vartheta(C_n) = \frac{n\cos(\pi/n)}{1 + \cos(\pi/n)}.$$

Since C_n has a node-transitive automorphism group, this implies that

$$\vartheta(\overline{C}_n) = 1 + \frac{1}{\cos(\pi/n)}.$$

♦

Example 6.1.12 (Kneser graphs) The Kneser graph K_k^n is defined on node set $\binom{[n]}{k}$, by connecting two k-sets if and only if they are disjoint ($1 \le k \le n$). Let us assume that $n \ge 2k$ to exclude the trivial case of a graph with no edges. The set of k-sets containing any fixed element of $[n]$ is stable, hence $\alpha(K_k^n) \ge \binom{n-1}{k-1}$. The Erdős–Ko–Rado Theorem asserts that this is the exact value; this fact will follow from our considerations below.

To compute the theta-function of this graph, we need the eigenvalues of its adjacency matrix. These are well known from coding theory:

$$(-1)^t \binom{n-k-t}{k-t}, \quad (t = 0, 1, \ldots, k).$$

The multiplicity of the eigenvalue with parameter t is $\binom{n}{t} - \binom{n}{t-1}$, but this is not important here. The largest eigenvalue in $\binom{n-k}{k}$ (the degree of each node), while the smallest is the next one, $-\binom{n-k-1}{k-1}$.

We apply the formula in Lemma 6.1.10, and get

$$\vartheta(K_k^n) \le \frac{n\binom{n-k-1}{k-1}}{\binom{n-k}{k} - \binom{n-k-1}{k-1}} = \binom{n-1}{k-1}. \tag{6.28}$$

Comparing with the lower bound on $\alpha(K_n^k)$, we see that

$$\vartheta(K_k^n) = \alpha(K_k^n) = \binom{n-1}{k-1}. \tag{6.29}$$

In particular, the Petersen graph K_2^5 has $\vartheta(K_2^5) = 4$. ♦

Example 6.1.13 (Paley graphs) The *Paley graph* Pal_p is defined for a prime $p \equiv 1 \pmod 4$. We take the $\{0, 1, \ldots, p-1\}$ as nodes, and connect two of them if their difference is a quadratic residue. It is clear that these graphs have a node-transitive automorphism group, and it is easy to see that they are self-complementary. So Corollary 6.1.7 applies, and gives that $\vartheta(\mathsf{Pal}_p) = \sqrt{p}$. To determine the stability number of Paley graphs is a difficult unsolved number-theoretic problem; it is conjectured that $\alpha(\mathsf{Pal}_p) = O((\log p)^2)$. ♦

Example 6.1.14 (Cycles with diagonals) For graphs with a node-transitive automorphism group, the existence of automorphism-invariant optima can be very useful. We illustrate this on the graph W_n obtained of an even cycle C_n with its longest diagonals added. In fact, we can restrict our attention to the case when $n = 4k$ is a multiple of 4, since otherwise W_n is bipartite and $\vartheta(G) = 2k$. We can observe the easy bounds $2k - 1 \le \vartheta(W_n) \le 2k$, since W_n has $2k - 1$ independent nodes and can be covered by $2k$ edges.

Let $V(W_n) = \{0, 1, \ldots, n-1\}$, where the nodes are labeled in the order of the original cycle. There exists an optimizing matrix A in (6.23) that is invariant under rotation, which means that it has only three different entries:

$$A_{uv} = \begin{cases} 1 + a, & \text{if } u - v \equiv \pm 1 \pmod{4k}, \\ 1 + b, & \text{if } u - v \equiv 2k \pmod{4k}, \\ 1, & \text{otherwise.} \end{cases}$$

It is easy to see that for every $(4k)$-th root of unity ε, the vector $(1, \varepsilon, \varepsilon^2, \dots \varepsilon^{n-1})$ is an eigenvector of A. (This is a complex vector, so if we want to stay in the real field, we have to consider its real and imaginary parts; but it is more convenient here to compute with complex vectors.) The eigenvalue λ_r corresponding to $\varepsilon = e^{2\pi i r/n}$ is easy to compute:

$$\lambda_0 = 4k + 2a + b,$$

and

$$\begin{aligned}\lambda_r &= 1 + a\varepsilon + \varepsilon^2 + \cdots + b\varepsilon^{2k} + \varepsilon^{2k+1} + \cdots + a\varepsilon^{n-1} = a(\varepsilon + \overline{\varepsilon}) + b\varepsilon^{2k} \\ &= 2a\cos\frac{r\pi}{2k} + b(-1)^r \qquad (r > 0).\end{aligned}$$

We want to minimize $\max_r \lambda_r$. Each $\lambda_r = \lambda_r(a, b)$ is a linear function of a and b, so we can find ϑ as the optimum of the linear program in 2 variables:

$$\vartheta(W_n) = \min_{a,b} \max_r \lambda_r(a, b). \tag{6.30}$$

There are many ways to do a back-of-the-envelope computation here; one gets that $a = -k$, $b = -k + k\cos(\pi/k)$ is an optimal solution, giving

$$\vartheta(W_n) = k + k\cos\frac{\pi}{k}. \tag{6.31}$$

The main point in this example is to illustrate that for graphs with a node-transitive automorphism group, the value of the theta function can be computed by a linear program, analogous to (6.30), where the number of unknowns is the number of orbits of the automorphism group on the edges. This may or may not lead to simple formula like in this case, but the computation is easy to perform, often even by hand. ♦

Example 6.1.15 (Self-polar polytopes) A polytope $P \subseteq \mathbb{R}^d$ is called *self-polar*, if $P^* = -P$. Note that this condition implies that for each vertex $\mathbf{v}$, the inequality $(-\mathbf{v})^{\mathsf{T}}\mathbf{x} \le 1$ defines a facet $F_{\mathbf{v}}$ of P, and we obtain all facets this way. We call two vertices $\mathbf{v}$ and $\mathbf{v}'$ of P *opposite*, if $\mathbf{v}'$ lies of $F_{\mathbf{v}}$. In other words, $\mathbf{v}^{\mathsf{T}}\mathbf{v}' = -1$, which shows that this is a symmetric relation. We call the polytope *strongly self-polar*, if it is inscribed in a ball centered at the origin, in other words, there is an $r > 0$ such that $|\mathbf{v}| = r$ for all vertices $\mathbf{v}$. For two opposite vertices we have $1 = \mathbf{v}^{\mathsf{T}}\mathbf{u} < |\mathbf{v}|\,|\mathbf{u}| = r^2$, and hence $r > 1$. It also follows that the distance of any facet from the origin is $1/r$, so the sphere with radius $1/r$ about the origin touches every facet.

In dimension 2, regular polygons with an odd number of vertices, with appropriate edge length, are strongly self-polar. It is known that for every dimension d and $\varepsilon > 0$ there exist strongly self-polar polytopes inscribed in a sphere with radius $r < 1 + \varepsilon$.

Let P be a strongly self-polar polytope in $\mathbb{R}^d$, and let G be the graph on $V(G) = V(P)$, in which two vertices are connected if and only of they are opposite. It is known that

$\chi(G) \geq d+1$. (The proof of this would lead us to a different area of advanced combinatorial methods, the use of algebraic topology.)

We can estimate $\vartheta(G)$ and $\vartheta(\overline{G})$ as follows. Let us label each vertex $\mathbf{v}$ of P with the vector

$$\mathbf{u_v} = \frac{1}{\sqrt{r^2+1}}\begin{pmatrix}\mathbf{v}\\ 1\end{pmatrix}.$$

This is trivially a unit vector in $\mathbb{R}^{d+1}$, and $\mathbf{u_v} \perp \mathbf{u_{v'}}$ for opposite vertices $\mathbf{v}$ and $\mathbf{v'}$. So $\mathbf{u}$ is an orthonormal representation of $\overline{G}$. Using the vector $\mathbf{e}_{d+1}$ as handle, we see that

$$\vartheta(\overline{G}) \leq \max_{\mathbf{v}} \frac{1}{(\mathbf{e}_{d+1}^{\mathsf{T}}\mathbf{u_v})^2} = r^2+1,$$

and

$$\vartheta(G) \geq \sum_{\mathbf{v}} (\mathbf{e}_{d+1}^{\mathsf{T}}\mathbf{u_v})^2 = \frac{n}{r^2+1}.$$

In particular, we see that $\chi(G)$ can be arbitrarily large while $\vartheta(\overline{G})$ can be arbitrarily close to 2. ⧫

6.2 Random graphs

It is a nontrivial problem to determine the theta function of a random graph. To start with a heuristic, recall Corollary 6.1.7: a self-complementary graph G with a node-transitive automorphism group has $\vartheta(G) = \sqrt{n}$.

Consider the most basic random graph $G(n, 1/2)$. This is, of course, not self-complementary, and its automorphism group is trivial, with high probability. However, its distribution is invariant under complementation and also under all permutations of the nodes. Informally, it is difficult to distinguish it from its complement (as it is difficult to distinguish any two random graphs with the same edge-density), and apart from a little variance in the degrees, it is difficult to distinguish any two nodes. So perhaps it does behave like a self-complementary graph with a node-transitive automorphism group would!

This heuristic predicts the right order of magnitude of $\vartheta(G(n, 1/2))$, namely it is of the order $\sqrt{n}$. However, no proof is known that would build on the heuristic above.

Theorem 6.2.1 *With high probability,*

$$\frac{1}{3}\sqrt{n} < \vartheta(G(n, \frac{1}{2})) < 3\sqrt{n}.$$

The method extends to estimating $\vartheta(G(n,p))$ (as usual, $G(n,p)$ denotes a random graph on n nodes with edge probability p): if p is a constant and $n \to \infty$, then with high probability,

$$\frac{1}{3}\sqrt{\frac{pn}{1-p}} < \vartheta(G(n,p)) < 3\sqrt{\frac{(1-p)n}{p}}.$$

Proof. First we prove the upper bound. Let A be the matrix defined by

$$A_{ij} = \begin{cases} -1, & \text{if } ij \in E, \\ 1, & \text{otherwise.} \end{cases}$$

Note that $\mathsf{E}(A_{ij}) = 0$ and $A_{ij}^2 = 1$ for $i \neq j$. The matrix A satisfies the conditions in Proposition 6.1.9, and hence we get that $\vartheta(G) \leq \lambda_{\max}(A)$, where $\lambda_{\max}(A)$ is the largest eigenvalue of A.

To estimate $\lambda_{\max}(A)$, we fix an integer $m = \Theta(n^{1/6})$, and note that $\lambda_{\max}(A)^{2m} \leq \operatorname{tr}(A^{2m})$.

Our next goal is to estimate the expectation of $\operatorname{tr}(A^{2m})$. This trace can be expressed as

$$\operatorname{tr}(A^{2m}) = \sum_{i_1,\ldots,i_{2m} \in V} A_{i_1 i_2} \ldots A_{i_{2m-1} i_{2m}} A_{i_{2m} i_1} = \sum_W A(W),$$

where $W = (i_1, \ldots, i_{2m}, i_1)$ is a closed walk on the complete graph (with loops) on n nodes. If we take expectation, then every term in which the defining walk traverses a non-loop edge an odd number of times gives 0. This implies that for the remaining terms, the number of different non-loop edges that are traversed is at most m, and since the graph traversed is connected, the number of different nodes that are touched is at most $m+1$. Let $V(W)$ denote the set of these nodes.

The main contribution comes from walks with $|V(W)| = m+1$; in this case no loop edges are used, the subgraph traversed is a tree, and every edge of it is traversed exactly twice. The traversing walk corresponds to selecting a starting node, drawing the tree in the plane as a rooted tree, and walking around it, keeping it on our right. It is well known that the number of different rooted plane trees with $m+1$ nodes is given by the Catalan number $\frac{1}{m+1}\binom{2m}{m}$. The walk determines a unique order of the $m+1$ nodes, corresponding to the order in which they are first encountered. We can associate these $m+1$ nodes with nodes in V in $n(n-1)\ldots(n-m)$ ways, and so the contribution of these terms is

$$T = \frac{1}{m+1}\binom{2m}{m} n(n-1)\ldots(n-m) \leq \frac{1}{m} 4^m n^{m+1} = \frac{n}{m}(4n)^m. \tag{6.32}$$

Those walks with $|V(W)| \leq m$ will contribute less, and we need a rather rough estimate of their contribution only. Consider a walk $W = (i_1, \ldots, i_{2m}, i_1)$ which covers a graph H on nodes $V(H) = V(W) = [r]$ that is not a tree. There is a non-loop edge (i_t, i_{t+1}) that is not a cut-edge of H. If it is traversed (at least) twice in the same direction, then we can write $W = W_1 + (i_t, i_{t+1}) + W_2 + (i_t, i_{t+1}) + W_3$, where W_1 is a walk from i_1 to i_t, W_1 is a walk from i_{t+1} to i_t, and W_3 is a walk from i_{t+1} to i_1 (Figure 6.2). Define a walk

$$W' = W_1 + (i_t, r+1, i_t) + \overleftarrow{W}_2 + W_3,$$

where $\overleftarrow{W}_2$ is the walk W_2 in reverse order.

If (i_t, i_{t+1}) is traversed once in each direction, then we have $W = W_1 + (i_t, i_{t+1}) + W_2 + (i_{t+1}, i_t) + W_3$, where W_1 is a walk from i_1 to i_t, W_2 is a walk from i_{t+1} to i_{t+1}, and W_3 is

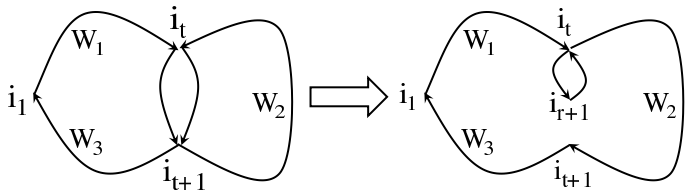

Figure 6.2

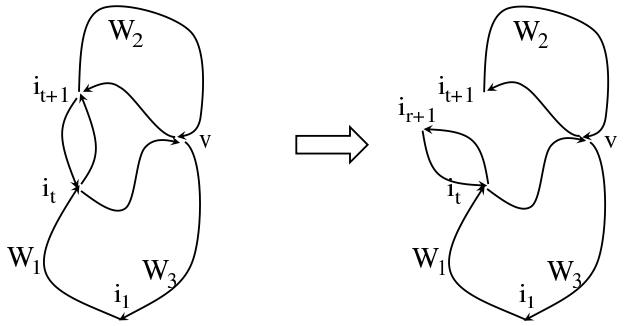

Figure 6.3

a walk from i_t to i_1 (Figure 6.3). Furthermore, the two closed walks $W_1 + W_3$ and W_2 must have a node v in common, since otherwise $i_t i_{t+1}$ would be a cut-edge of H. Reversing the order on the segment of W between these two copies of v, we get a new walk W_0 that has the same contribution, and which passes $i_t i_{t+1}$ twice in the same direction, and so W_0' can be defined as above. Let us define $W' = W_0'$.

In both cases, the new closed walk W' covers $r + 1$ nodes. Furthermore, by simple computation,

$$\mathsf{E}(A(W)) \leq 2\mathsf{E}(A(W')).$$

Every closed walk covering $[r + 1]$ arises at most $4m^3$ ways in the form of W': we can get W back by identifying $r + 1$ with one of the other nodes ($r \leq m$ choices), and then flipping the

segment between two occurrences of the same node (fewer than $4m^2$ choices). Hence

$$\begin{aligned}\sum_{|V(W)|=r} \mathsf{E}(A(W)) &= \binom{n}{r} \sum_{V(W)=[r]} \mathsf{E}(A(W)) \le 2\binom{n}{r} \sum_{V(W)=[r]} \mathsf{E}(A(W')) \\ &\le 8m^3 \binom{n}{r} \sum_{V(W)=[r+1]} \mathsf{E}(A(W)) \\ &= 8m^3 \binom{n}{r}\binom{n}{r+1}^{-1} \sum_{|V(W)|=r+1} \mathsf{E}(A(W)) \\ &= \frac{8(r+1)m^3(m+1)}{n-r} \sum_{|V(W)|=r+1} \mathsf{E}(A(W)) \\ &\le \frac{1}{2} \sum_{|V(W)|=r+1} \mathsf{E}(A(W)),\end{aligned}$$

by the choice of m, if n is large enough. This implies that

$$\mathsf{E}(\mathrm{tr}(A^{2m})) = \sum_W \mathsf{E}(A(W)) \le \sum_{j=0}^{m} \frac{T}{2^j} < 2T.$$

Hence by Markov's Inequality and (6.32),

$$\begin{aligned}\mathsf{P}\Big(\lambda_{\max}(A) \ge 3\sqrt{n}\Big) &= \mathsf{P}\Big(\lambda_{\max}(A)^{2m} \ge (9n)^m\Big) \le \mathsf{P}\Big(\mathrm{tr}(A^{2m}) \ge (9n)^m\Big) \\ &\le \frac{\mathsf{E}(\mathrm{tr}(A^{2m}))}{(9n)^m} \le \frac{2T}{(9n)^m} = o(1).\end{aligned}$$

This proves that with high probability, $\lambda_{\max}(A) < 3\sqrt{n}$ as claimed.

To prove the lower bound, it suffices to invoke Lemma 6.1.4:

$$\vartheta(G(n,\frac{1}{2})) = \vartheta(\overline{G}(n,\frac{1}{2})) \ge \frac{n}{\vartheta(G(n,\frac{1}{2}))} > \frac{n}{3\sqrt{n}} = \frac{1}{3}\sqrt{n}$$

with high probability. □

6.3 Polyhedral combinatorics of the theta function

6.3.1 Polarity and antiblocking

Polarity. Let K be a convex body containing the origin as an interior point. The *polar* of K is defined by

$$K^* = \{\mathbf{x} \in \mathbb{R}^d : \ \mathbf{x}^\mathsf{T}\mathbf{y} \le 1 \ \forall \mathbf{y} \in P\}.$$

It is clear that K^* is a convex body as well, containing the origin in its interior. For every convex body K we have $(K^*)^* = K$.

The polar of a polytope (containing the origin) is a polytope. For every vertex $\mathbf{v}$ of P, the inequality $\mathbf{v}^\mathsf{T}\mathbf{x} \leq 1$ defines a facet of P^*, and vice versa. The vector $\mathbf{v}$ is a normal vector of the facet $\mathbf{v}^\mathsf{T}\mathbf{x} \leq 1$. More generally, if $\mathbf{v}_0, \ldots, \mathbf{v}_m$ are the vertices of a k-dimensional face F of P, then

$$F^\perp = \{\mathbf{x} \in P^* : \ \mathbf{v}_0^\mathsf{T}\mathbf{x} = 1, \ldots, \mathbf{v}_m^\mathsf{T}\mathbf{x} = 1\}$$

defines a $d - k - 1$-dimensional face of P^*. Furthermore, $(F^\perp)^\perp = F$.

For every convex cone C, we can form its *polar cone* C^*, defined by

$$C^* = \{x \in \mathbb{R}^n : x^\mathsf{T} y \geq 0 \ \forall y \in C\}.$$

This is again a convex cone. If C is closed (in the topological sense), then we have $(C^*)^* = C$.

The polar cone of the cone of semidefinite matrices (within the space of symmetric $n \times n$ matrices) is itself.

Antiblocking. A closed convex set K is called a *convex corner*, if $K \subseteq \mathbb{R}^d_+$ and whenever $\mathbf{x} \in K$, $\mathbf{y} \in \mathbb{R}^d$ and $0 \leq \mathbf{y} \leq \mathbf{x}$ then $\mathbf{y} \in K$. A polytope which is a convex corner will be called a *corner polytope*. The *antiblocker* of a convex corner K is defined by

$$K^{\mathrm{abl}} = \{\mathbf{x} \in \mathbb{R}^d_+ : \ \mathbf{x}^\mathsf{T}\mathbf{y} \leq 1 \forall \mathbf{y} \in K\}$$

(see Figure 6.4).

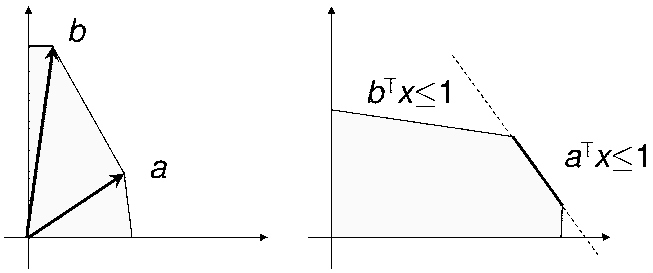

Figure 6.4: A 2-dimensional corner polytope and its antiblocker

The antiblocker of a convex corner is a convex corner. The antiblocker of a corner polytope is a corner polytope. For every convex corner K, we have $(K^{\mathrm{abl}})^{\mathrm{abl}} = K$.

The correspondence between faces of a corner polytope P and P^{abl} is again a bit more complicated than for the polars. The nonnegativity constraints $x_i \geq 0$ always define facets, and they don't correspond to vertices in the antiblocker. All other facets of P correspond to vertices of P^{abl}. Not every vertex of P defines a facet in P^{abl}. The origin is a trivial exceptional vertex, but there may be further exceptional vertices. We call a vertex $\mathbf{v}$ *dominated*, if there is another vertex $\mathbf{w}$ such that $\mathbf{v} \leq \mathbf{w}$. Now a vertex of P defines a facet of P^* if and only if it is not dominated.

6.3.2 The stable set polytope

Stable sets and cliques give rise to important polyhedra associated with graphs. After summarizing those basic properties of these polyhedra that we need, we show that orthogonal representations provide an interesting related convex body, with nice duality properties.

The *stable set polytope* $\mathrm{STAB}(G)$ of a graph G is the convex hull of incidence vectors of all stable sets. This gives us a polytope in $\mathbb{R}^V$. The stability number $\alpha(G)$ can be obtained by maximizing the linear function $\sum_{i\in V} x_i$ over this polytope, which suggests that methods from linear programming can be used here.

With this goal in mind, we have to find a system of linear inequalities whose solution set is exactly the polytope $\mathrm{STAB}(G)$. It would be best to find a minimal such system, which is unique. If we can find this system, then the task of computing the stability number $\alpha(G)$ of G reduces to maximizing $\sum_{i\in V} x_i$ subject to these constraints, which means solving a linear program. Unfortunately, this system of linear inequalities is in general exponentially large and very complicated. But if we find at least some linear inequalities valid for the stable set polytope, then solving the linear program we get an upper bound on $\alpha(G)$, and for special graphs, we get the exact value.

So we want to find linear inequalities (constraints) valid for the incidence vector of every stable set. We start with the trivial *non-negativity constraints*:

$$x_i \geq 0 \quad (i \in V). \tag{6.33}$$

The fact that the set is stable is reflected by the *edge constraints*:

$$x_i + x_j \leq 1 \quad (ij \in E). \tag{6.34}$$

Inequalities (6.33) and (6.34) define the *fractional stable set polytope* $\mathrm{FSTAB}(G)$. Integral points in $\mathrm{FSTAB}(G)$ are exactly the incidence vectors of stable sets, but $\mathrm{FSTAB}(G)$ may have other (non-integral) vertices, and is in general larger than $\mathrm{STAB}(G)$. The case of equality has a nice characterization.

Proposition 6.3.1 $\mathrm{STAB}(G) = \mathrm{FSTAB}(G)$ *if and only if* G *is bipartite.* □

Let $\alpha^f(G)$ denote the maximum of $\sum_i x_i$ over $\mathbf{x} \in \mathrm{FSTAB}(G)$. Trivially $\alpha(G) \leq \alpha^f(G)$, and $\alpha^f(G)$ is computable in polynomial time (since (6.33) and (6.34) describe a linear program defining α^f).

We can strengthen the edge constraints if the graph has larger cliques. Every clique B gives rise to a *clique constraint*:

$$\sum_{i\in B} x_i \leq 1. \tag{6.35}$$

Inequalities (6.33) and (6.35) define a polytope $\mathrm{QSTAB}(G)$, the *clique-constrained fractional stable set polytope* of G. Since cliques in G correspond to stable sets in $\overline{G}$ and vice versa, it is easy to see that $\mathrm{QSTAB}(G)$ is just the antiblocker of $\mathrm{STAB}(\overline{G})$.

Again, we can introduce a corresponding relaxation of the stability number, namely the quantity $\alpha^*(G)$ defined as the maximum of $\sum_i x_i$ over $\mathbf{x} \in \mathrm{QSTAB}(G)$. This quantity is a sharper upper bound on $\alpha(G)$ then α^f, but it is NP-hard to compute.

The polytope $\mathrm{QSTAB}(G)$ is contained in $\mathrm{FSTAB}(G)$, but is still larger than $\mathrm{STAB}(G)$ in general. The case of equality leads us to an interesting and rich class of graphs, of which we give a very brief survey.

6.3.3 Perfect graphs

A graph G is called *perfect*, if for every induced subgraph G' of G, we have $\omega(G') = \chi(G')$. Every bipartite graph is perfect, since they satisfy $\omega(G) = \chi(G) = 2$ (if they have an edge) or $\omega(G) = \chi(G) = 1$ (if they have no edge), and their induced subgraphs are also bipartite. Figure 6.5 shows some perfect and non-perfect graphs.

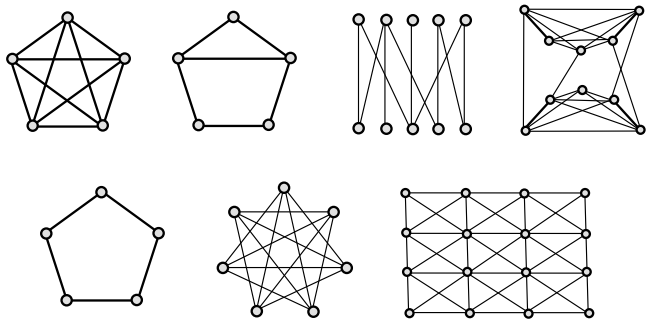

Figure 6.5: Some perfect graphs (first row) and some non-perfect graphs (second row).

To be perfect is a rather strong structural property; nevertheless, many interesting classes of graphs are perfect (bipartite graphs, their complements and their linegraphs, interval graphs, comparability and incomparability graphs of posets, chordal graphs, split graphs, etc.).

The following deep characterization of perfect graphs was conjectured by Berge in 1961 and proved by Chudnovsky, Robertson, Seymour and Thomas in 2006.

Theorem 6.3.2 (The Strong Perfect Graph Theorem) *A graph is perfect if and only if neither the graph nor its complement contains a chordless odd cycle longer than* 3. □

As a corollary we can state the "The Weak Perfect Graph Theorem":

Theorem 6.3.3 *The complement of a perfect graph is perfect.* □

From this theorem it follows that in the definition of perfect graphs we could replace the equation $\omega(G') = \chi(G')$ by $\alpha(G') = \chi(\overline{G'})$. Perfectness can also be characterized in terms of the stable set polytope:

Theorem 6.3.4 $\mathrm{STAB}(G) = \mathrm{QSTAB}(G)$ *if and only if* G *is perfect.* □

Based on our remark above, the condition $\mathrm{STAB}(G) = \mathrm{QSTAB}(G)$ is equivalent to saying that $\mathrm{STAB}(\overline{G})$ and $\mathrm{STAB}(G)$ are antiblockers, which is a condition symmetric in G and $\overline{G}$. So Theorem 6.3.4 implies Theorem 6.3.3.

6.3.4 Orthogonality constraints

For every orthonormal representation $(\mathbf{u}_i, \mathbf{c})$ of G, we consider the linear constraint

$$\sum_{i \in V} (\mathbf{c}^{\mathsf{T}} \mathbf{u}_i)^2 x_i \leq 1, \tag{6.36}$$

which we call an *orthogonality constraint*. The solution set of non-negativity and orthogonality constraints is denoted by $\mathrm{TSTAB}(G)$. It is clear that TSTAB is a closed, full-dimensional, convex set. The orthogonality constraints are valid if x is the indicator vector of a stable set of nodes (cf. (6.1)), and therefore they are valid for $\mathrm{STAB}(G)$. Furthermore, every clique constraint is an orthogonality constraint. Indeed, for every clique B, the constraint $\sum_{i \in B} x_i \leq 1$ is obtained from the orthogonal representation

$$i \mapsto \begin{cases} \mathbf{e}_1, & i \in B, \\ \mathbf{e}_i, & \text{otherwise}, \end{cases} \qquad \mathbf{c} = \mathbf{e}_1.$$

Hence

$$\mathrm{STAB}(G) \subseteq \mathrm{TSTAB}(G) \subseteq \mathrm{QSTAB}(G) \tag{6.37}$$

for every graph G.

There are several other characterizations of TSTAB. These are based on an extension of the theta-function to the case when we are also given a weighting $w : V \to \mathbb{R}_+$. Generalizing the formulas in Remark 6.1.3, the quantity $\vartheta(G, w)$ can be defined by any of the following formulas:

$$\vartheta(G, w) = \min\Big\{ \max_{i \in V} \frac{w_i}{(\mathbf{c}^{\mathsf{T}} \mathbf{u}_i)^2} : \ \mathbf{u} \text{ ONR of } G,\ |\mathbf{c}| = 1 \Big\} \tag{6.38}$$

$$= \min\Big\{ t \geq 2 : |\mathbf{y}_i|^2 = t - w_i,\ \mathbf{y}_i^{\mathsf{T}} \mathbf{y}_j = -\sqrt{w_i w_j}\ (ij \in \overline{E}) \Big\} \tag{6.39}$$

$$= \min\big\{ \max_{i \in V} (Y_{ii} + w_i) : \ Y \succeq 0,\ Y_{ij} = -\sqrt{w_i w_j}\ (ij \in \overline{E}) \big\} \tag{6.40}$$

$$= \max\Big\{ \sum_{i,j \in V} w_i w_j Z_{ij} : \ Z \succeq 0,\ Z_{ij} = 0\ (ij \in E),\ \sum_i Z_{ii} = 1 \Big\} \tag{6.41}$$

$$= \max\Big\{ \sum_{i \in V} w_i (\mathbf{d}^{\mathsf{T}} \mathbf{v}_i)^2 : \ \mathbf{v} \text{ ONR of } \overline{G},\ |\mathbf{d}| = 1 \Big\}. \tag{6.42}$$

The equivalence of (6.38)–(6.42) can be obtained extending the proof Theorem 6.1.2 to the node-weighted version (at the cost of a little more computation).

For every orthonormal representation $\mathbf{u} = (\mathbf{u}_i : i \in V)$ with handle $\mathbf{c}$, we call the vector $((\mathbf{c}^\mathsf{T}\mathbf{u}_i)^2 : i \in V)$ the *profile* of the orthogonal representation. We can state two further characterizations of TSTAB(G):

Proposition 6.3.5 (a) $\mathbf{x} \in \mathrm{TSTAB}(G)$ *if and only if* $\vartheta(\overline{G}, \mathbf{x}) \leq 1$.

(b) *The body* TSTAB(G) *is exactly the set of profiles of dual orthonormal representations of* G.

Proof. (a) follows from (6.42).

(b) The profile of every dual orthonormal representation belongs to TSTAB(G); this is equivalent to (6.8). Conversely, let $\mathbf{x} \in \mathrm{TSTAB}(G)$. Then $\vartheta(\overline{G}, \mathbf{x}) \leq 1$ by (a), so (6.38) implies that there is a dual orthonormal representation $\mathbf{v}$ of G with handle $\mathbf{d}$ for which $x_i \leq \mathbf{d}^\mathsf{T}\mathbf{v}_i$ for all nodes $i \in V$. Thus the vectors $\mathbf{v}'_i = (x_i/\mathbf{d}^\mathsf{T}\mathbf{v}_i)\mathbf{v}_i$ satisfy $\mathbf{d}^\mathsf{T}\mathbf{v}'_i = x_i$. The vectors $\mathbf{v}'_i$ are not of unit length, but the vectors

$$\mathbf{v}''_i = \begin{pmatrix} \mathbf{v}'_i \\ \sqrt{1 - |\mathbf{v}'_i|^2}\mathbf{e}_i \end{pmatrix} \quad \text{and} \quad \mathbf{d}'' = \begin{pmatrix} \mathbf{d} \\ 0 \end{pmatrix}$$

form a dual orthonormal representation of G with profile $\mathbf{x}$. □

The last characterization of TSTAB(G) is equivalent to the following duality result.

Corollary 6.3.6 TSTAB($\overline{G}$) *is the antiblocker of* TSTAB(G). □

Before stating the next theorem, recall that a *vertex* of a convex body K is a boundary point $\mathbf{v}$ that is the unique point of intersection of all hyperplanes supporting K at $\mathbf{v}$. This means that there is a pointed convex cone containing K with $\mathbf{v}$ as its vertex. This is to be distinguished from an *extreme point*, which is the unique point of intersection of a hyperplane supporting K with K.

Theorem 6.3.7 *The vertices of* TSTAB(G) *are exactly the incidence vectors of stable sets in* G.

Does this imply that TSTAB(G) = STAB(G)? Of course not, since TSTAB(G) (as every convex body) is the convex hull of its extreme points, but not necessarily of its vertices.

Proof. The vector $\mathbb{1}_A$, where A is a stable set of nodes, is the unique common point of the supporting hyperplanes $x_i = 1$ ($i \in A$) and $x_i = 0$ ($i \in V \setminus A$), and so it is a vertex of TSTAB(G).

Conversely, let $\mathbf{z} = (z_i : i \in V)$ be a vertex of TSTAB(G). If $z_i = 0$ for some node i, then we can delete i: We get a graph G' for which $\mathrm{TSTAB}(G') = \mathrm{TSTAB}(G) \cap \{z_i = 0\}$, and so $\mathbf{z}\,|_{V\setminus\{i\}}$ is a vertex of TSTAB(G'), and we can proceed by induction.

So we may assume that $z_i > 0$ for all $i \in V$. Since $\mathbf{z} \in \mathrm{TSTAB}(G)$, we can write $z_i = (\mathbf{d}^\mathsf{T}\mathbf{v}_i)^2$ for some dual orthonormal representation $(\mathbf{v}_i : i \in V)$ of G and unit vector $\mathbf{d}$.

Let $\mathbf{a}^\mathsf{T}\mathbf{x} \leq 1$ be a hyperplane that supports TSTAB(G) at $\mathbf{z}$. Then $\mathbf{a} \in \text{TSTAB}(\overline{G})$ by Corollary 6.3.6, and hence there is an orthonormal representation $(\mathbf{u}_i : i \in V)$ of G and unit vector $\mathbf{c}$ such that $a_i = (\mathbf{c}^\mathsf{T}\mathbf{u}_i)^2$. By the same argument as in the derivation of (6.17), we get

$$\mathbf{d} = \sum_i (\mathbf{c}^\mathsf{T}\mathbf{u}_i)^2 (\mathbf{d}^\mathsf{T}\mathbf{v}_i)\mathbf{v}_i.$$

Multiplying by any vector $\mathbf{y}$, we get

$$\mathbf{d}^\mathsf{T}\mathbf{y} = \sum_i (\mathbf{c}^\mathsf{T}\mathbf{u}_i)^2 (\mathbf{d}^\mathsf{T}\mathbf{v}_i)(\mathbf{v}_i^\mathsf{T}\mathbf{y}) = \sum_i a_i (\mathbf{d}^\mathsf{T}\mathbf{v}_i)(\mathbf{v}_i^\mathsf{T}\mathbf{y}).$$

Thus if $\mathbf{y}$ is not orthogonal to $\mathbf{c}$, then the point $\mathbf{z}'$ defined by $z'_i = (\mathbf{d}^\mathsf{T}\mathbf{v}_i)(\mathbf{v}_i^\mathsf{T}\mathbf{y})/(\mathbf{d}^\mathsf{T}\mathbf{y})$ is contained in the supporting hyperplane $\sum_i a_i x_i = 1$. This holds for every supporting hyperplane at $\mathbf{z}$. Since $\mathbf{z}$ is a vertex, the only common point of hyperplanes supporting TSTAB(G) at $\mathbf{z}$ is $\mathbf{z}$ itself (here we use that $\mathbf{z}$ is a vertex, not just an extreme point). Thus $\mathbf{z}' = \mathbf{z}$, which means that $\mathbf{v}_i^\mathsf{T}\mathbf{y} = (\mathbf{d}^\mathsf{T}\mathbf{v}_i)(\mathbf{d}^\mathsf{T}\mathbf{y})$ for all i (here we use that $z_i \neq 0$ for all i). This linear equation must also hold for the vectors $\mathbf{y}$ orthogonal to $\mathbf{c}$. So we get that $\mathbf{v}_i = (\mathbf{d}^\mathsf{T}\mathbf{v}_i)\mathbf{d}$, and since $\mathbf{v}_i$ is a unit vector, we get $\mathbf{v}_i = \mathbf{d}$. So no two vectors $\mathbf{v}_i$ are orthogonal, and thus G has no edges. So $\mathbf{z} = \mathbb{1}_V$ is the incidence vector of a stable set as claimed. □

Example 6.3.8 Consider the graph C_5, with node set $\{1, \ldots, 5\}$. The polytope STAB(C_5) has 11 vertices (the origin, the basic unit vectors, and the incidence vectors of non-adjacent pairs of nodes). The facets are defined by the non-negativity constraints, edge constraints, and single further inequality

$$x_1 + x_2 + x_3 + x_4 + x_5 \leq 2. \tag{6.43}$$

Since C_5 has no triangles, we have QSTAB = FSTAB. This polytope has a single vertex $(\frac{1}{2}, \ldots, \frac{1}{2})$ in addition to the incidence vectors of stable sets.

Turning to TSTAB, we know by Theorem 6.3.7 that it has 11 vertices, the same vertices as STAB. The umbrella construction we have used before gives a point

$$\Big(\frac{1}{\sqrt{5}}, \ldots, \frac{1}{\sqrt{5}}\Big)^\mathsf{T} \in \text{TSTAB}(C_5) \tag{6.44}$$

which is not in STAB(G) by (6.43). Applying the umbrella construction to the complement, and scaling, we get an orthogonality constraint

$$x_1 + \cdots + x_5 \leq \sqrt{5}, \tag{6.45}$$

showing that the special vertex of FSTAB(C_5) does not belong to TSTAB(C_5). ♦

This example shows that not every orthogonality constraint follows from the clique constraints. In fact, the number of essential orthogonality constraints is infinite unless the graph is perfect.

Proposition 6.3.9 TSTAB(G) *is polyhedral if and only if the graph is perfect.*

Proof. If G is perfect, then STAB(G) = QSTAB(G) by Theorem 6.3.4, and (6.37) implies that TSTAB(G) = STAB(G) = QSTAB(G), so TSTAB(G) is polyhedral. To prove the converse, suppose that TSTAB(G) is polyhedral, then Theorem 6.3.7 implies that TSTAB(G) = STAB(G). We can apply this argument to $\overline{G}$, since TSTAB($\overline{G}$) = TSTAB(G)$^{\text{abl}}$ is also polyhedral; we get that

$$\text{TSTAB}(G) = \text{TSTAB}(\overline{G})^{\text{abl}} = \text{STAB}(\overline{G})^{\text{abl}} = \text{QSTAB}(G).$$

So STAB(G) = TSTAB(G) = QSTAB(G), which implies that G is perfect by Theorem 6.3.4. □

6.4 Applications

6.4.1 Shannon capacity

In the introduction, we have described how to use orthogonal representations to determine the Shannon zero-error capacity of the pentagon. What happens with other confusion graphs?

Let V be an alphabet with confusion graph $G = (V, E)$. To describe the confusion graph of longer messages, we use the strong product of two graphs. In these terms, $\alpha(G^{\boxtimes k})$ is the maximum number of non-confusable words of length k: words composed of elements of V, so that for every two words there is at least one i ($1 \leq i \leq k$) such that the i-th letters are different and non-adjacent in G, i.e., non-confusable. It is easy to see that

$$\alpha(G \boxtimes H) \geq \alpha(G)\alpha(H). \tag{6.46}$$

This implies that

$$\alpha(G^{\boxtimes(k+l)}) \geq \alpha(G^{\boxtimes k})\alpha(G^{\boxtimes l}), \tag{6.47}$$

and hence

$$\alpha(G^{\boxtimes k}) \geq \alpha(G)^k. \tag{6.48}$$

The *Shannon capacity* (zero-error capacity, if we want to be pedantic) of a graph G is the value

$$\Theta(G) = \lim_{k\to\infty} \alpha(G^{\boxtimes k})^{1/k}. \tag{6.49}$$

Inequality (6.47) implies, via Fekete's Lemma, that the limit exists, and (6.48) implies that

$$\Theta(G) \geq \alpha(G). \tag{6.50}$$

Rather little is known about this graph parameter for general graphs. For example, it is not known whether $\Theta(G)$ can be computed for all graphs by any algorithm (polynomial or not), although there are several special classes of graphs for which this is not hard. The behavior of $\Theta(G)$ and the convergence in (6.49) are rather erratic.

Let us describe a few facts we do know. First, let us generalize the argument from the Introduction bounding $\Theta(C_4)$. Let $\overline{\chi}(G)$ denote the minimum number of complete subgraphs covering the nodes of G (this is the same as the chromatic number of the complementary graph.) Trivially

$$\alpha(G) \le \overline{\chi}(G). \tag{6.51}$$

Any covering of G by $\overline{\chi}(G)$ cliques and of H by $\overline{\chi}(H)$ cliques gives a "product covering" of $G \boxtimes H$ by $\overline{\chi}(G)\overline{\chi}(H)$ cliques, and so

$$\overline{\chi}(G \boxtimes H) \le \overline{\chi}(G)\overline{\chi}(H). \tag{6.52}$$

Hence

$$\alpha(G^{\boxtimes k}) \le \overline{\chi}(G^{\boxtimes k}) \le \overline{\chi}(G)^k,$$

and thus

$$\Theta(G) \le \overline{\chi}(G). \tag{6.53}$$

It follows that if $\alpha(G) = \overline{\chi}(G)$, then $\Theta(G) = \alpha(G)$; for such graphs, nothing better can be done than reducing the alphabet to the largest mutually non-confusable subset. In particular, this answers the Shannon capacity problem for perfect graphs.

Instead of $\overline{\chi}$, we can use ϑ to bound the Shannon capacity:

$$\alpha(G^k) \le \vartheta(G^k) \le \vartheta(G)^k,$$

which implies

Proposition 6.4.1 *For every graph G,*

$$\Theta(G) \le \vartheta(G).$$

Since $\vartheta(C_5) = \sqrt{5}$, we get that $\Theta(C_5) = \sqrt{5}$. This argument can be generalized to an infinite class of graphs:

Corollary 6.4.2 *If G is a self-complementary graph with a node-transitive automorphism group, then $\Theta(G) = \sqrt{n}$.*

Proof. The diagonal in $G \boxtimes \overline{G}$ is stable, so $\alpha(G \boxtimes G) = \alpha(G \boxtimes \overline{G}) \ge n$, and hence $\Theta(G) \ge \sqrt{n}$. On the other hand, $\Theta(G) \le \vartheta(G) = \sqrt{n}$ by Corollary 6.1.7. □

Example 6.4.3 (Paley graphs II) Paley graphs form a class of graphs to which this corollary applies, and whose Shannon capacity can be determined exactly: $\Theta(\mathrm{Pal}_p) = \vartheta(\mathrm{Pal}_p) = \sqrt{p}$. Assuming that the stability number of a Paley graph is polylogarithmic in p (as conjectured), for this infinite family of graphs the Shannon capacity is much higher than the stability number. ♦

The tensor product construction in the proof of Theorem 6.1.8 shows that if G has an orthonormal representation in dimension c, and H has an orthonormal representation in dimension d, the $G \boxtimes H$ has an orthonormal representation in dimension cd. It follows that the minimum dimension of any orthonormal representation is an upper bound on $\Theta(G)$. It is not hard to show that this bound is never better than $\vartheta(G)$. However, if we consider orthogonal representations over fields of finite characteristic, then the analogue of ϑ is not defined, but the minimum dimension provides a better bound on the Shannon capacity than ϑ in some cases.

We'll return to the Shannon capacity in a quantum communication setting in last section of this chapter.

6.4.2 Coloring perfect graphs

Perhaps the most important consequence of the formulas proved in Section 6.1 is that the value of $\vartheta(G)$ is polynomial time computable. More precisely,

Theorem 6.4.4 *There is an algorithm that computes, for every graph G and every $\varepsilon > 0$, a real number t such that*

$$|\vartheta(G) - t| < \varepsilon.$$

The running time of the algorithm is polynomial in n and $\log(1/\varepsilon)$.

Algorithms proving this theorem can be based on almost any of our formulas for ϑ. The simplest is to refer to Theorem 6.1.2 giving a formulation of $\vartheta(G)$ as the optimum of a semidefinite program (6.3), and the polynomial time solvability of semidefinite programs.

The significance of this fact is underlined if we combine it with Theorem 6.1.1: The two important graph parameters $\alpha(G)$ and $\chi(\overline{G})$ are both NP-hard, but they have a polynomial time computable quantity sandwiched between them. This fact in particular implies:

Corollary 6.4.5 *The stability number and the chromatic number of a perfect graph are polynomial time computable.*

Using the algorithms of Corollary 6.4.5 one can compute more than just these values: one can compute a maximum stable set and an optimal coloring in a perfect graph in polynomial time.

Theorem 6.4.4 extends to the weighted version of the theta function. Maximizing a linear function over STAB(G) or QSTAB(G) is NP-hard; but, surprisingly, TSTAB behaves much better: *Every linear objective function can be maximized over* TSTAB(G) *(with an arbitrarily small error) in polynomial time.* This applies in particular to $\vartheta(G)$, which is the maximum of $\sum_i x_i$ over TSTAB(G).

6.5 Orthogonal representations and quantum physics

We conclude our treatment of orthogonal representations with sketching some applications of them in quantum physics. As it happens, the applications we are going to discuss are all about entanglement, so let us introduce this notion first.

As the basic setup in quantum physics, the state of a physical system can be described by a vector of unit norm in a (complex) Hilbert space. In the simple systems we need, this space will be finite dimensional, so it may be considered as $\mathbb{C}^d$ for an appropriate nonnegative integer d. We write $\mathbf{u} \cdot \mathbf{v}$ for the inner product of two vectors $\mathbf{u}, \mathbf{v} \in \mathbb{C}^d$, to emphasize that it is used in the Hilbert space sense: $\mathbf{u} \cdot \mathbf{v} = \sum_i \overline{u}_i v_i$.

A *measurement* on a quantum state $\mathbf{u}$ is performed by applying a self-adjoint linear operator A to it. The simplest measurement is the orthogonal projection of the state vector onto a one-dimensional subspace. If $\mathbf{a}$ is the vector generating this subspace (which is unique up to a scalar of absolute value 1), then the projection is $(\mathbf{u} \cdot \mathbf{a})\mathbf{a}$. Such an experiment can be thought of as checking a particular property of $\mathbf{u}$, and the probability that the property checks out is $|\mathbf{a} \cdot \mathbf{u}|^2$.

The fact that we assign vectors to discrete objects like properties suggests analogies with orthogonal representations (and perhaps other representations treated in this book). As it turns out, this is more than just an analogy; we describe three special problems, where quantum physics makes strong use of the theory of orthogonal representations. It is quite possible that other connections can be found.

6.5.1 Capacity of quantum channels

The most successful area of applying quantum physics in computer science, at least so far, is quantum information theory. Using phenomena of quantum physics (for example, entanglement of particles), one can create communication channels that are more efficient than classical communication channels.

Consider two quantum physical systems A and B. Separately, the their states can be described by unit vectors $\mathbf{x}$ and $\mathbf{y}$ in a complex Hilbert space $\mathbb{C}^d$. (For these applications, we may restrict our attention to finite dimensional state spaces, and we simplify notation by assuming that both systems have a d-dimensional state space.) The state of the union of the two systems can be described by a vector in the tensor product $\mathbb{C}^d \otimes \mathbb{C}^d$. If the two systems

in states $\mathbf{x}$ and $\mathbf{y}$ are "independent" (unentangled), their joint state is $\mathbf{x} \circ \mathbf{y}$. However, this will not be the case in general, as a vector $\mathbf{z} \in \mathbb{C}^d \otimes \mathbb{C}^d$ cannot be written as the tensor product of two vectors in $\mathbb{C}^d$. In this case, the state $\mathbf{z}$ is called *entangled*. Such a state can be written as

$$\mathbf{z} = \sum_{u,v=1}^{d} S_{uv}(\mathbf{e}_u \circ \mathbf{e}_v),$$

where $\{\mathbf{e}_1, \ldots, \mathbf{e}_d\}$ is the standard basis in $\mathbb{C}^d$, and the coefficients S_{uv} are complex numbers with $\sum_{u,v} |S_{uv}|^2 = 1$. Sometimes it is convenient to describe $\mathbf{z}$ by the $d \times d$ complex matrix $S = (S_{ij})$. A *maximally entangled* state is

$$\mathbf{z} = \sum_{i=1}^{d} \frac{1}{\sqrt{d}} \mathbf{e}_i \circ \mathbf{e}_i.$$

Entanglement leads to rather paradoxical behavior of particles; this was pointed out by Einstein, Podolsky and Rosen in 1935. Consider a pair of maximally entangled particles with a 2-dimensional state space: say two electrons, whose spin can be either "up" or "down". Their maximally entangled state is $\frac{1}{\sqrt{2}}\mathbf{e}_1 \circ \mathbf{e}_1 + \frac{1}{\sqrt{2}}\mathbf{e}_2 \circ \mathbf{e}_2$. Such a pair is often called an Einstein–Podolsky–Rosen pair or EPR pair.

Suppose that Alice and Bob split an EPR pair between themselves (while it remains in the same entangled state), and they travel to different far away places. If Alice measures the state of her electron, she will find it in one of the states $\mathbf{e}_1$ and $\mathbf{e}_2$ with the same probability. Say it is in state e_1, then the entangled state collapses to $\mathbf{e}_1 \circ \mathbf{e}_1$ immediately. This is long-range action faster than light, so it contradicts special relativity—but only almost. If Bob measures the state of his electron, then it will be in state $\mathbf{e}_1$; but if they do the measurements in different order what they observe will come out the same. You can play around with more EPR pairs to convince yourself that no information can be transmitted between Alice and Bob.

But such a strange behavior can be utilized, in particular in information theory and computer science. In the example above, Alice and Bob do obtain *the same* random bit, which is very secure, since by the principles of quantum physics, nobody can learn this bit without destroying the entanglement.

One can ask for analogues of the Shannon capacity in quantum information theory. It turns out that ϑ provides an upper bound for the quantum physical version Shannon capacity just like it does for classical channels.

Example 6.5.1 While entanglement does not allow to send information directly (faster than the speed of light) between two people, it does allow to make better use of classical communication channels.

We need a graph G with the following properties: (a) its node set can be partitioned into q complete d-subgraphs $H_1, \ldots, H_q$; (b) there are further complete d-subgraphs $H_{d+1}, \ldots, H_r$

such that every edge of G belongs to one of $H_1, \ldots, H_r$; (c) $\alpha(G) < q$; (d) G has a dual orthonormal representation $\mathbf{v}$ in $\mathbb{R}^d$. Several constructions of such graphs are known, it is perhaps easiest to describe such a graph with $d = 4$ and $q = 6$. The node set of this graph can be defined as the set of vectors in $\mathbb{R}^4$ with coordinates $0, 1$ or -1, where the number of nonzero coordinates is $1, 2$ or 4, and of vectors that are negatives of each other we only keep one. We connect two of these by an edge if they are orthogonal. By definition, this graph has a dual orthonormal representation in $\mathbb{R}^4$. To verify that properties (a)–(d) above are satisfied is left to the reader.

Next we construct a (noisy) classical communication channel. Its input alphabet is V, its output alphabet is $[r]$, and on input $a \in V$ it outputs one of the indices i for which $a \in V(H_i)$. The output i is chosen randomly and uniformly from all such i. Thus two inputs a and b are confusable if and only if there is a chance that they lead to the same output, i.e., if there is an i such that $a, b \in V(H_i)$. This is equivalent to $ab \in E$, so the confusability graph of the channel is G. The number of inputs that can be used without any danger of confusion is $\alpha(G) < q$.

We use the dual orthonormal representation to show that in the presence of entanglement, one can safely transmit q different one-letter messages. We construct two copies a quantum system with a d-dimensional state space $\mathbb{C}^d$, and prepare them by bringing them to the maximally entangled state

$$\mathbf{u} = \sum_{i=1}^{d} \frac{1}{\sqrt{d}} \mathbf{e}_i \circ \mathbf{e}_i.$$

Alice and Bob get one of the two copies each. Now if Alice wants to send message m to Bob ($1 \leq m \leq q$), then she measures the state of her side in the orthonormal basis $\{\mathbf{v}_s : s \in V(H_m)\}$. The result of the measurement will be a random element $\mathbf{v}_s$ of this basis. After that, the entangled system will be in the state $\mathbf{v}_s \circ \mathbf{v}_s$.

Alice sends s to Bob through the noisy channel. Bob receives an index i for which $s \in V(H_i)$. Since H_i is complete, Bob can measure the state of his side in the orthonormal basis $\{\mathbf{v}_t : t \in V(H_i)\}$, and determine s. Since there is a unique m, $1 \leq m \leq q$, for which $s \in V(H_m)$, he determines m. ♦

So we see that using an entangled state can improve the zero-error capacity of a noisy channel. How much is the improvement? It turns out that the theta function is still an upper bound. One can use an entanglement-assisted channel repeatedly to gain in its capacity just like in the classical case. Since the theta function is multiplicative, it remains an upper bound on the zero-error capacity.

Theorem 6.5.2 *Suppose that Alice and Bob are connected by a classical noisy channel, where the input alphabet has confusability graph G. In addition, there is an entangled state $\mathbf{u} \in \mathbb{C}^d \circ \mathbb{C}^d$, where Alice has access to the first factor and Bob has access to the second*

factor. Then the maximum number of one-letter messages that Alice can transmit without the possibility of confusion is bounded by $\vartheta(G)$.

Proof. Let the noisy channel have input alphabet X and output alphabet Z. For each input x, let $Z_x \subseteq Z$ be the set of outputs that occur with positive probability. (Since we are interested in zero-error, the actual probabilities don't matter.) Two elements x and y are confusable if and only if $Z_x \cap Z_y \neq \emptyset$.

The other part of the equipment, the shared entangled state can be described by a vector in $\mathbb{R}^d \otimes \mathbb{R}^d$, which we consider as a $d \times d$ complex matrix S such that $S \cdot S = \mathrm{tr}(S\overline{S}^{\mathsf{T}}) = \sum_{u,v} |S_{uv}|^2 = 1$.

If Alice wants to transmit a message $i \in [m]$, then she performs a measurement on her half of $\mathbf{u}$ depending on i, and transmits a message $x \in X$ that depends on i and the result of the measurement. The measurement i leading to message x is an orthogonal projection A_i^x onto a subspace of $\mathbb{C}^d$. The fact that always exactly one message $x \in X$ must be sent means, by the laws of quantum physics, that $A_i^x A_i^y = 0$ for $x \neq y$ and $\sum_i A_i^x = I$ for each $x \in X$. This measurement collapses $\mathbf{u}$ to the state $\mathbf{u}'$ described by the matrix $A_i^x S$.

Once Bob gets the message (more exactly, he gets some message $z \in Z_x$), he performs a measurement on $\mathbf{u}'$ (depending on z), which has m possible outcomes; the outcome of this measurement is supposed to coincide with Alice's intended message i.

Bob's measurements consist of m orthogonal projections $B_1^z, \ldots, B_m^z$ onto appropriate subspaces of $\mathbb{C}^d$. Similarly as before, we have $B_i^z B_j^z = 0$ for $i \neq j$ and $\sum_i B_i^z = I$. Measurement B_j^z collapses the state to $A_i^x S B_j^z$. The fact that Bob is able to recover i means that $A_i^x S B_j^z = 0$ if $z \in Z_x$ and $i \neq j$.

We claim that

$$(A_i^x S) \cdot (A_j^y S) = 0 \tag{6.54}$$

in the following cases: (a) $i = j,\ x \neq y$; (b) $i \neq j,\ x = y$; (c) $i \neq j,\ xy \in E$. Case (a) is easy, since we can write $(A_i^x S) \cdot (A_i^y S) = \mathrm{tr}(\overline{S}^{\mathsf{T}} A_i^x A_i^y S)$, and, as we have seen, $A_i^x A_i^y = 0$. In cases (b) and (c) there is an element $z \in Z_x \cap Z_y$, which gives

$$(A_i^x S) \cdot (A_j^y S) = \mathrm{tr}(A_j^y S\overline{S}^{\mathsf{T}} A_i^x) = \mathrm{tr}\big(A_j^y S\big(\sum_k B_k^z\big)\overline{S}^{\mathsf{T}} A_i^x\big) = \sum_k \mathrm{tr}\big(A_j^y S B_k^z \overline{S}^{\mathsf{T}} A_i^x\big) = 0,$$

since for every k, either $A_j^y S B_k^z = 0$ or $B_k^z \overline{S}^{\mathsf{T}} \overline{A}_i^x = 0$. This proves (6.54).

Define an $X \times X$ matrix M by

$$M_{x,y} = \sum_{i,j=1}^{m} (A_i^x S) \cdot (A_j^y S).$$

If $xy \in E$, then $M_{x,y} = 0$, since every term is zero by conditions (a) and (c) above. It is easy

to see that M is positive semidefinite. Furthermore, using (b),

$$\mathrm{tr}(M) = \sum_x \sum_{i,j=1}^{m} (A_i^x S) \cdot (A_j^x S) = \sum_x S \cdot S = m,$$

and

$$\mathrm{tr}(JM) = \sum_{x,y,i,j} \mathrm{tr}\big(\overline{S}^{\mathsf{T}} A_i^x A_j^y S\big) = \sum_{i,j} \mathrm{tr}\Big(\overline{S}^{\mathsf{T}} \big(\sum_x A_i^x\big)\big(\sum_y A_j^y\big) S\Big) = m^2 \mathrm{tr}(\overline{S}^{\mathsf{T}} S) = m^2.$$

This shows that the matrix $\frac{1}{m}M$ almost fulfils definition 6.12 of ϑ, except that it is not necessarily real. But $\frac{1}{2m}(M + \overline{M}^{\mathsf{T}})$ is real, and has the same trace and sum of entries, so $\vartheta(G) \geq m$. □

One can consider quantum physical communication channels more general than entanglement-assisted classical channels. The theta-function can be generalized to such channels, so that it remains an upper bound on the zero–error capacity of the channel.

6.5.2 Contextuality

We describe a connection between the theta function and a rather fundamental issue in quantum physics. Consider a quantum system, and let $e_1, \ldots, e_n$ be observable events. Construct a graph G on $V = [n]$ in which $ij \in E$ if and only if e_i and e_j are exclusive (cannot occur simultaneously). We call G the *exclusivity graph* of the system $e_1, \ldots, e_n$ of events.

We start the system in a state $\mathbf{u}$, and observe an event e_i. As we know, this observation changes the state, so we cannot observe all of the other events. But we assume that we can restart from the same state $\mathbf{u}$, and then we can observe other events. Repeating this many times, we can find the probability P_i that event e_i occurs.

What can we say about the sum $P_1 + \cdots + P_n$? In the classical setting, when $e_1, \ldots, e_n$ are observable events in a probability space, the sum $P_1 + \cdots + P_n$ would be the expected number of events that occur simultaneously. It is trivial that this number is at most $\alpha(G)$, so we get the inequality

$$P_1 + \cdots + P_n \leq \alpha(G). \tag{6.55}$$

The same inequality can be derived in quantum physics, if we assume its "hidden variable" interpretation. This interpretation arose from the objection to the nondeterministic–random interpretation of quantum events: it suggests that if we knew the exact state of each particle (its "hidden parameters"), then we could predict quantum events with certainty. Another consequence would be that even if we are not able to observe certain events simultaneously, they do actually occur. In the well known parable of Schrödinger's cat, the cat in the closed chamber at a given moment in time is either dead or alive—we just don't know which. Similarly, if for each experiment a definite subset of the events e_i occurs (of which we only

know the outcome of one), then the calculation of the expected number of events that occur goes just like above, and we get (6.55).

Bell was the first to suggest that "Bell Inequalities" related to (6.55) could be experimentally verified (or rather falsified); since their derivation depends on the theory of hidden variables, disproving such an inequality would disprove the hidden variable theory (at least in its simplest form).

From basic quantum physical principles (not using hidden variables) one can only derive the weaker inequality

$$P_1 + \cdots + P_n \leq \vartheta(G). \tag{6.56}$$

Indeed, state $\mathbf{u}$ is a unit vector in a complex Hilbert space H. Observing the event e_i means projecting $\mathbf{u}$ to a one-dimensional subspace, spanned by a unit vector $\mathbf{a}_i \in H$; the probability that the event occurs is just the squared length of the projection. Thus $P_i = |\mathbf{u} \cdot \mathbf{a}_i|^2$, and

$$P_1 + \cdots + P_n = \sum_i |\mathbf{u} \cdot \mathbf{a}_i|^2.$$

Now two events that exclude each other must correspond to projections onto orthogonal vectors, and hence $\mathbf{a}$ is a dual orthogonal representation of G in a Hilbert space. This is a complex Hilbert space, but it is not hard to see that the value of the maximum in definition (6.5) of the theta function does not change if we take complex Hilbert space instead of the real. Hence we get

$$P_1 + \cdots + P_n = \sum_i |\mathbf{u} \cdot \mathbf{a}_i|^2 \leq \vartheta(G).$$

From even weaker (simpler) principles, one gets an even weaker inequality. The *exclusivity principle* says that the sum of probabilities of mutually exclusive events is at most 1. This implies the inequality

$$\sum_{i \in B} P_i \leq 1$$

for every clique B in G; in the language of Section 6.3.2, the vector $(P_1, \ldots, P_n)$ must belong to the clique-constrained fractional stable set polytope QSTAB(G), and so it follows that

$$P_1 + \cdots + P_n \leq \alpha^*(G). \tag{6.57}$$

Two special cases have been studied extensively. In the Clauser–Horne–Shimony–Holt version of the Bell experiment, one creates two maximally entangled electrons, whose spins can be both "up" or both "down". These are sent in opposite directions to two far away observers, called (as usual) Alice and Bob. Alice has two possible settings a and a' of her equipment measuring the spin, and Bob has similarly two settings b and b'. Let (a_+b_+) denote

the event that Alice measures in setting a, Bob in setting b, and they both measure "up". We have 16 analogous events. Repeating the experiment many times, they can determine the probabilities of these events.

The *Clauser–Horne–Shimony–Holt–Bell inequality* is a linear inequality concerning these probabilities, which can be derived from the assumption of hidden variables. Since there are trivial linear equations relating these probabilities, this inequality can be written in several equivalent forms. For us, the following form will be relevant:

$$P(a_+b_+) + P(a_-b_-) + P(a'_+b_+) + P(a'_-b_-) + P(a_+b'_+) + P(a_-b'_-) + P(a'_+b'_-) + P(a'_-b'_+) \leq 3 \tag{6.58}$$

(note the twist in the last two terms!). This is a special case of (6.55), applied to the exclusivity graph G of the 8 events $(a_+b_+), (a_-b_-), (a'_+b_+), (a'_-b_-), (a_+b'_+), (a_-b'_-), (a'_+b'_-), (a'_-b'_+)$. This graph is depicted in Figure 6.6(b); looking at the picture, it is easy to see that $\alpha(G) = 3$. Furthermore, we have $\vartheta(G) = 2 + \sqrt{2} \approx 3.414\ldots$ (see Example 6.1.14), so the inequality

$$P(a_+b_+) + P(a_-b_-) + P(a'_+b_+) + P(a'_-b_-) + P(a_+b'_+) + P(a_-b'_-) + P(a'_+b'_-) + P(a'_-b'_+) \leq 2 + \sqrt{2} \tag{6.59}$$

follows by (6.56) without the hypothesis of hidden variables, just by the laws of quantum physics.

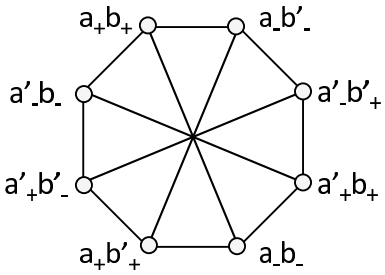

Figure 6.6: The exclusivity graph V_8 of the Clauser–Horne–Shimony–Holt experiment. The edge between (say) $a'_-b'_+$ and $a_-b'_-$ indicates that they cannot occur simultaneously: for these two events, test b' returns different results.

The Klyachko-Can-Binicioğlu-Shumovsky experiment works in a similar spirit, and leads to a more complicated measurement for a simpler graph, namely C_5.

Which of these bounds is the "truth"? After a long line of increasingly sophisticated experiments, recent reports by several groups claim to have eliminated all the implicit assumptions ("loopholes"), and show that the bound (6.58) does not hold in general. This can be considered as a disproof of the "hidden variable" interpretation of quantum physics.

It can also be shown that (6.56) is best possible in the sense that for every simple graph G one can construct mathematical models of quantum physical systems in which equality holds.