A leisurely traveling salesman tour through several facets of cuts

Monique Laurent

CWI & Tilburg University

Foreword

Beste Jan Karel,

Dit is een kort verhaal over sneden (cuts) in grafen, als afscheidskado en dank voor al je werk en inzet voor het CWI. Hopelijk heb jij voortaan meer tijd en kunnen we vaker van je inzicht over allerlei 'traveling salesman' vraagstukken genieten!

Veel dank en ik wens je veel plezier in je nieuwe leven na de CWI directie.

Prologue

Cuts in graphs are very basic objects. However, they often illustrate in a beautiful way several recent developments in combinatorial optimization, and they have fascinating connections to other mathematical areas like metric theory, distance geometry, combinatorial matrix theory. My co-authored book *Geometry of Cuts and Metrics* which just appeared when I joined CWI in 1997, covers some of them. [As an aside, this earned me my first experience with the generous and thoughtful work environment at CWI!]

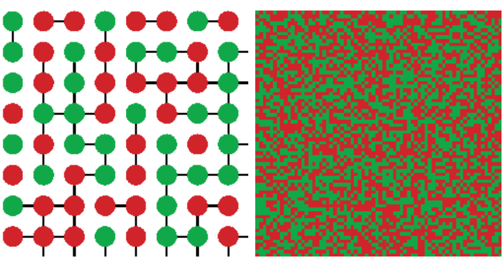

Cuts have numerous applications, for example, to cluster analysis (cluster a set of data into groups of most similar objects), to statistical physics (minimize the energy of spin glass systems), or to VLSI design (minimize the number of 'vias' needed to route nets on several layers in chip design). My focus here will be to travel from a bird-eye perspective through some chosen mathematical questions and results where cuts play a central role. This is not a scientific paper, there is no pretention in completeness, and I will mostly omit to say precisely who has done what.

Complexity of max-cut

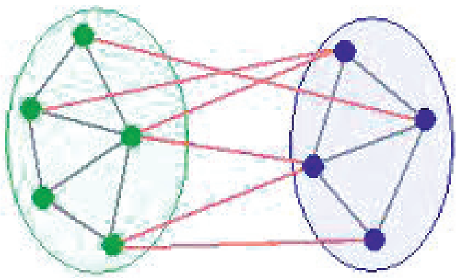

Given a graph $G = (V = [n], E)$, a *cut* is the set of edges that are 'cut' by some partition of the nodes into two groups. Say $(S, V \setminus S)$ is such a partition, then the corresponding cut consists of all edges $ij \in E$ with $i \in S$ and $j \notin S$ (or conversely).

If weights $w \in \mathbb{R}^E$ are assigned to the edges of G, then the *max-cut problem* asks to find a cut whose weight $\sum_{i \in S, j \notin S} w_{ij}$ is maximum. Two remarks are in order about the complexity of max-cut.

First, if all weights are non-negative and, instead of a maximum cut, one wants a *minimum weight non-empty cut*, then the problem is polynomial time solvable, using maximum flow techniques and Fulkerson's celebrated max-flow min-cut theorem.

On the other hand, max-cut is NP-hard, even if all weights are equal to 1 and the graph is cubic. Karp [1972] proposed the following reduction of max-cut from the partition problem: Given integers $a_1, \ldots, a_n \in \mathbb{N}$, define the edge weights $w_{ij} = a_i a_j \ \forall i, j$. Then the weight of a cut is $(\sum_{i \in S} a_i)(\sum_{i \notin S} a_i) \leq (\sum_{i \in V} a_i)^2/4$, with equality if and only if $\sum_{i \in S} a_i = (\sum_{i \in V} a_i)/2$, i.e., the sequence $a_1, \ldots, a_n$ can be partitioned.

However, max-cut can be solved in polynomial time over the class of planar graphs; working with the planar dual graph, the problem can indeed be reduced to shortest path combined with maximum weight matching computations. Polynomial time solvability extends to the larger class of graphs which do not contain a K_5 minor.

Cut polyhedra

Define the *cut polytope* $\mathrm{CUT}_n^{\square}$ as the convex hull of the incidence vectors of all cuts in the complete graph K_n. Thus solving max-cut amounts to linear optimization over the cut polytope. Define analogously the *cut cone* CUT_n as the conic hull of the incidence vectors of all cuts. A beautiful symmetry property is that all the facets of the cut polytope can derived by a simple 'switching' operation from the facets of the cut cone (where 'switching' means changing signs along a cut). Hence it suffices to find the facets passing through the origin... Well, as max-cut is NP-hard, this remains probably a hopeless task! Nevertheless many classes of facets are known. In particular, facets arise within the class of hypermetric inequalities mentioned in (2) below. It is known that the subclass of triangle inequalities (3) (in fact, their projections) suffices to describe the cut cone for graphs with no K_5 minor. Gaining good information about the facets of cut polyhedra is important for branch-and-cut type algorithms for max-cut, but it is also relevant to other fields, including ℓ_1-metrics, or the study of Bell inequalities in quantum information theory, as mentioned below.

Cuts and ℓ_1-metrics

Any cut corresponds to a very simple distance on V: if two points of V are cut they are at distance 1 and otherwise they are at distance 0. Call a distance d ℓ_1-*realizable* if there exist vectors $u_1, \ldots, u_n$ in some space $\mathbb{R}^k$ satisfying $d_{ij} = \|u_i - v_j\|_1 \quad \forall i, j$. Clearly, any cut corresponds to an ℓ_1-realizable distance (then the u_i's can even be chosen within $\{0, 1\}$) and one can show that the distances that are ℓ_1-realizable are precisely the members of the cut cone CUT_n. Thus, the hardness of max-cut implies that testing ℓ_1-realizability is also a hard problem.

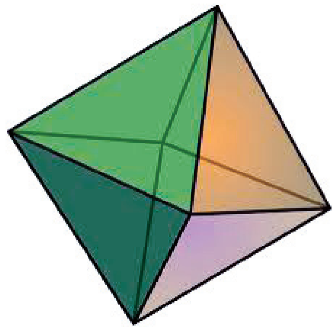

The octahedron: The unit ball of the ℓ_1-norm in 3D

In comparison, testing ℓ_2-realizability can be done in polynomial time - using semidefinite programming. Indeed, there exist vectors $u_1, \ldots, u_n \in \mathbb{R}^k$ $(k \geq 1)$ such that $\|u_i - u_j\|_2 = d_{ij} \quad \forall i, j$ if and only if there exists a positive semidefinite matrix X satisfying the linear conditions:

$$X_{ii} + X_{jj} - 2X_{ij} = d_{ij}^2 \; \forall i, j \in [n]. \tag{1}$$

Nevertheless the above characterization of ℓ_1-metrics as members of the cut cone implies that any valid inequality for the cut cone gives a necessary condition for ℓ_1-realizability. Among the many known classes of valid inequalities, the class of *hypermetric inequalities* is of particular interest. Given integers $b_1, \ldots, b_n \in \mathbb{Z}$ with $\sum_i b_i = 1$, the corresponding hypermetric inequality is

$$\sum_{1 \leq i < j \leq n} b_i b_j d_{ij} \leq 0. \tag{2}$$

For instance, for $b_i = b_j = 1 = -b_k$ and all other entries equal to 0, we get the *triangle inequalities*, claiming that d is a metric:

$$d_{ij} \leq d_{ik} + d_{jk} \tag{3}$$

(thus motivating the name *hyper*metric). For $b = (1, 1, 1, -1, -1, 0, \ldots, 0)$, we get the *pentagonal inequality*:

$$d_{12} + d_{13} + d_{23} + d_{45} - \sum_{i \in \{1,2,3\}, j \in \{4,5\}} d_{ij} \leq 0. \tag{4}$$

Although there are infinitely many hypermetric inequalities, the cone they define can yet be shown to be a polyhedral cone. This result is based on a deep connection between hypermetrics and geometry of numbers (Voronoi and Delaunay polytopes in lattices).

Semidefinite approximation for max-cut

Assume all edge weights are non-negative. Here is a simple random algorithm to get a good cut: Assign the nodes independently, with probability 1/2, to either side of the cut. Then an edge is cut with probability 1/2 and one gets a cut whose expected weight is half the total weight of edges, and thus at least half max-cut.

How to do better?

Goemans and Williamson [1995] introduced an ingenious construction for constructing a better approximate cut, based on semidefinite programming combined with a clever 'random hyperplane rounding' technique:

1. Find unit vectors $v_1, \dots, v_n \in \mathbb{R}^n$ maximizing $\sum_{ij \in E} w_{ij}(1 - v_i^T v_j)$.
2. Pick a random hyperplane, say with normal r.
3. This hyperplane induces a partition of the space, and thus a cut, depending on the sign of $r^T v_i$.

The beautiful property of this construction is that the expected weight of the obtained cut is at least 0.878 times max-cut.

The key fact in the analysis is that the probability that edge ij is cut depends only on the angle between vectors v_i, v_j, namely it is equal to $\arccos(v_i^T v_j)/\pi$. Hence the expected weight of the cut is equal to

$$\sum_{ij \in E} w_{ij} \frac{\arccos(v_i^T v_j)}{\pi} = \underbrace{\sum_{ij \in E} w_{ij} \frac{1 - v_i^T v_j}{2}}_{:=\text{SDP}} \underbrace{\frac{2}{\pi} \frac{\arccos v_i^T v_j}{1 - v_i^T v_j}}_{\geq \alpha_{\text{GW}}} \geq \text{ max-cut } \alpha_{\text{GW}},$$

where $\alpha_{\text{GW}} := \min_{\vartheta \in [0,\pi]} \frac{2}{\pi} \frac{\vartheta}{1 - \cos \vartheta} \sim 0.878.$

Indeed, the value SDP is at least max-cut. To see it, note that one can formulate max-cut as the following integer quadratic problem:

$$\text{max-cut} = \max \sum_{ij \in E} w_{ij}(1 - x_i x_j)/2 \;\text{ s.t. }\; x_1, \dots, x_n \in \{\pm 1\}. \tag{5}$$

Relaxing the condition: $x_i \in \{\pm 1\}$ by the condition: v_i lies in the unit sphere S^{n-1}, we get the following 'vector relaxation' of max-cut:

$$\text{SDP} := \max \sum_{ij \in E} w_{ij}(1 - v_i^T v_j)/2 \;\text{ s.t. }\; v_1, \dots, v_n \in S^{n-1}. \tag{6}$$

Letting $X = (v_i^T v_j)$ denote the Gram matrix of the v_i's, (6) can be reformulated as the following *semidefinite program:*

$$\text{SDP} = \max \sum_{ij \in E} w_{ij}(1 - X_{ij})/2 \ \text{ s.t. } X_{ii} = 1 \ \forall i, \ X \succeq 0. \tag{7}$$

Thus the value SDP, together with optimal vectors v_i's for (6), can be computed in polynomial time (to arbitrary precision) using, e.g., interior point algorithms.

The Goemans-Williamson approximation $\alpha_{\text{GW}} \sim 0.878$ improves the trivial 1/2 approximation. As of today, this remains the best known polynomial-time approximation guarantee for max-cut. Whether the GW approximation can be improved is intimately related to the *Unique Games conjecture* of Khot [2002] in complexity theory. Namely, if the Unique Games conjecture holds, then the Goemans-Williamson approximation is best possible.

A trigonometric puzzle about angles of vectors

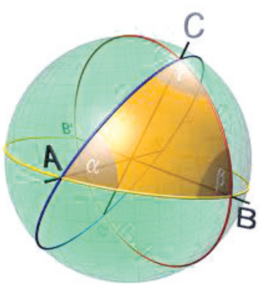

A spherical triangle:
$\alpha + \beta + \gamma \leq 2\pi$

Let $v_1, \ldots, v_n$ be given vectors on the unit sphere S^{n-1}. A natural question is what are the linear conditions satisfied by their angles ϑ_{ij} (= $\arccos v_i^T v_j$).

A basic well known geometric fact in 3D is that the pairwise angles of three vectors v_1, v_2, v_3 satisfy:

$$\vartheta_{12} \leq \vartheta_{13} + \vartheta_{23}, \tag{8}$$

$$\vartheta_{13} \leq \vartheta_{12} + \vartheta_{23},$$

$$\vartheta_{23} \leq \vartheta_{12} + \vartheta_{13},$$

$$\vartheta_{12} + \vartheta_{13} + \vartheta_{23} \leq 2\pi. \tag{9}$$

But, what about the pairwise angles of five vectors ?
Is it true that they satisfy:

$$\vartheta_{12} + \vartheta_{13} + \vartheta_{23} + \vartheta_{45} \leq \sum_{i \in \{1,2,3\}, j \in \{4,5\}} \vartheta_{ij} \tag{10}$$

as well as

$$\sum_{1 \leq i < j \leq 5} \vartheta_{ij} \leq 6\pi \ ? \tag{11}$$

Well, the answer is **yes** and the above 'random hyperplane rounding' argument permits to show this.

Indeed, say we have an inequality $\sum_{ij} w_{ij} d_{ij} \leq w_0$, valid for $\text{CUT}_n^{\square}$, and a collection of unit vectors $v_1, \ldots, v_n$. As in steps 2-3 of the Goemans-Williamson procedure described above, pick a random hyperplane and consider the corresponding random cut, whose expected weight is equal to $\sum_{ij} w_{ij}\vartheta_{ij}/\pi$, and is at most the maximum value w_0. This shows that $\sum_{ij} w_{ij}\vartheta_{ij} \leq w_0\pi$ holds for the pairwise angles of unit vectors. As an illustration, relation (8) corresponds to the triangle inequality (3), relation (10) corresponds to the pentagonal inequality (4), and (11) corresponds to the switched inequality, claiming that a cut in K_5 has at most six edges.

Spectrahedra and the elliptope

While the feasible region of a linear program is a polyhedron, the term *spectrahedron* has been coined for the feasible region of a semidefinite program. In particular, the feasible region of the semidefinite relaxation (7) of max-cut is a spectrahedron, called the *elliptope* and denoted as $\mathcal{E}_n$.

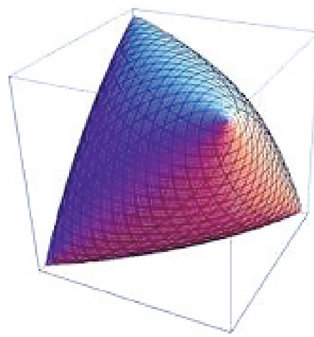

Fig. 1: The elliptope $\mathcal{E}_3$

The elliptope $\mathcal{E}_n$ consists of all $n \times n$ positive semidefinite matrices with diagonal entries one. Such matrices correspond to correlation matrices in statistics and partly for this reason their geometric properties have been much studied. For instance, it is known that a correlation matrix which is an extreme point of the elliptope $\mathcal{E}_n$ has rank r satisfying $\binom{r+1}{2} \leq n$. It is also known that the elliptope, although it is a non-polyhedral convex set, does have polyhedral faces, some of them being inherited from the cut polytope. The possible dimensions for the polyhedral faces of $\mathcal{E}_n$ have been characterized.

By construction, the elliptope $\mathcal{E}_n$ is a convex relaxation of the cut polytope $\mathrm{CUT}_n^{\square}$. For $n = 3$, there are four cuts in K_3 and the cut polytope $\mathrm{CUT}_3^{\square}$ is a simplex in $\mathbb{R}^3$. Fig. 1 shows the elliptope $\mathcal{E}_3$. Notice that it has four 'corners', corresponding to the four cuts. Also, $\mathcal{E}_3$ can be seen as the geometric shape obtained by 'inflating' the cut polytope while keeping its skeleton rigid (which earns $\mathcal{E}_3$ its alternative name: 'the pillow').

A basic natural question is to understand what are the characteristic properties of spectrahedra. By definition, a spectrahedron K is the solution set of a semidefinite program. That is, K can be written as the set of solutions $x \in \mathbb{R}^n$ of a so-called *linear matrix inequality (LMI)*:

$$A(x) := A_0 + x_1 A_1 + \ldots + x_n A_n \succeq 0. \tag{12}$$

Clearly, any spectrahedron K must be convex, and a basic closed semi-algebraic set, which means that K can be described by finitely many polynomial inequalities (indeed, simply write that the principal minors of the matrix $A(x)$ must be non-negative). A deep algebraic result of Helton and Vinnikov [2006] characterizes spectrahedra in $\mathbb{R}^2$ (in terms of convex rigidity), but a characterization is not known in dimension $n \geq 3$.

The next figure shows another spectrahedron: the set $\mathcal{T}_4$ of all Toeplitz matrices (i.e., with identical entries on each diagonal) in the elliptope $\mathcal{E}_4$, together with its dual body. Alternatively, $\mathcal{T}_4$ is the convex hull of the trigonometric curve $(\cos(t), \cos(2t), \cos(3t))$ for $t \in \mathbb{R}$.

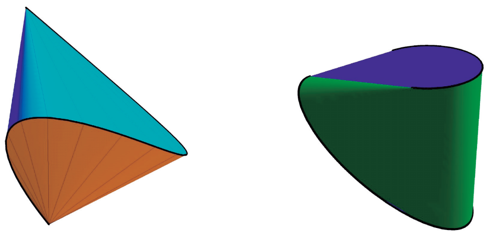

For the purpose of optimization one is also interested in working with projections of spectrahedra, which corresponds to lifting the problem in higher dimension. Call a set K *SDP representable* if it can be obtained as the projection of a spectrahedron. That is, there is an LMI:

$$A(x,z) = A_0 + x_1 A_1 + \ldots + x_n A_n + z_1 B_1 + \ldots + z_m B_m$$

such that $x \in K$ if and only if $A(x,z) \succeq 0$ for some $z \in \mathbb{R}^m$. Necessarily K must be convex and semi-algebraic, but again no full characterization is known. Characterizing SDP representable sets is a major open problem in the field of semidefinite programming and convex algebraic optimization.

As an illustration, the unit ball in $\mathbb{R}^n$ of the ℓ_2-norm is a spectrahedron, as

$$\left\{ x \in \mathbb{R}^n \mid \sum_{i=1}^{n} x_i^2 \leq 1 \right\} = \left\{ x \in \mathbb{R}^n \mid \begin{pmatrix} 1 & x^T \\ x & I_n \end{pmatrix} \succeq 0 \right\}.$$

In optimization terms, optimizing over the second order cone, also known as the Lorentz cone or the ice-cream cone: $\left\{ (x_0, x) \in \mathbb{R}^{n+1} \mid x_0 \geq \sqrt{\sum_{i=1}^{n} x_i^2} \right\}$, can be cast as an instance of semidefinite programming. On the other hand, the unit ball in $\mathbb{R}^2$ of the ℓ_4-norm:

$$K = \{(x,y) \in \mathbb{R}^2 \mid x^4 + y^4 \leq 1\},$$

is not a spectrahedron (since it is not rigidly convex), but it is SDP representable. Indeed K can be realized as the projection of a spectrahedron by adding two more variables. Namely, $(x,y) \in K$ if and only if there exists $(u,v) \in \mathbb{R}^2$ satisfying $x^2 \leq u$, $y^2 \leq v$, $u^2 + v^2 \leq 1$, which can equivalently be formulated using the LMI's:

$$\begin{pmatrix} 1 & x \\ x & u \end{pmatrix} \succeq 0, \ \begin{pmatrix} 1 & y \\ y & v \end{pmatrix} \succeq 0, \ \begin{pmatrix} 1 & u & v \\ u & 1 & 0 \\ v & 0 & 1 \end{pmatrix} \succeq 0.$$

Distance geometry and matrix completions

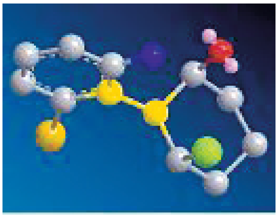

The *distance geometry problem* can be stated as follows: Given a graph $G = (V, E)$ and a (partial) set of distances $d = (d_{ij})_{ij \in E}$, find (if it exists) a realization of d in $\mathbb{R}^k$, i.e., find vectors $u_1, \ldots, u_n \in \mathbb{R}^k$ such that $d_{ij} = \|u_i - u_j\|_2$ for all $ij \in E$. This contains, for instance, the molecular conformation problem in molecular chemistry, where one wants to reconstruct the 3D structure of a molecule from the knowledge of certain interatomic distances, or sensor network localization problems.

If the dimension k is fixed and part of the input, then the problem is NP-hard (already for $k = 1$, where there is an easy reduction from the partition problem). However, the exact complexity of deciding existence of an ℓ_2-realization in any dimension $k \geq 1$ is not known. Note that this amounts to deciding whether the semidefinite program:

$$X_{ii} + X_{jj} - 2X_{ij} = d_{ij}^2 \ (ij \in E), \ \ X \succeq 0$$

(already encountered in (1) above) has a feasible solution. However, the exact complexity of deciding feasibility of a semidefinite program is not known. Let us make a brief excursion to the complexity of semidefinite programming.

Consider the problem of testing feasibility of the LMI (12), where $A_1, \ldots, A_n$ are given integer symmetric matrices. Clearly, this problem belongs to NP *in the real number model of computation*, since one can test whether a matrix is positive semidefinite in polynomial time - using Gaussian elimination. However, the problem is not known to be in NP *in the bit number model of complexity* (although Ramana [1997] could show that it belongs to NP if and only if it belongs to co-NP - using a clever extended duality theory). It is indeed not clear how to come up with a short certificate. For instance, the following semidefinite program:

$$\begin{pmatrix} 2 & x \\ x & 1 \end{pmatrix} \succeq 0, \ \begin{pmatrix} x & 2 \\ 2 & 2x \end{pmatrix} \succeq 0$$

admits a unique *irrational* solution: $x = \sqrt{2}$. Deciding the complexity status of testing feasibility of a semidefinite program is a major open problem. On the positive side, Khachiyan and Porkolab [1997] have shown that one can test existence of a rational solution of the LMI (12) in polynomial time when *fixing the number n of variables*, thus extending the result of Lenstra [1983] about polynomial time solvability of integer LP in fixed dimension to SDP.

On the positive side, the distance geometry problem can be solved in polynomial time for the class of chordal graphs, and for the class of graphs with no K_4 minor. The latter result relies on the following observation. Clearly, if d has an ℓ_2-realization then it can be completed to a metric $\hat{d}$ on V; moreover one can test in polynomial time whether such a metric completion exists. It turns out that existence of a metric completion is also sufficient to ensure existence of an ℓ_2-realization in the case when the graph has no K_4 minor. More strongly, if d has an ℓ_2-realization, then it has a metric completion $\hat{d}$ which satisfies *all* valid inequalities for the cut cone CUT_n. (In sophisticated terms, ℓ_2-realizability implies ℓ_1-realizability.)

Two examples of distances with no ℓ_2-realization are shown in Fig. 2. As a first example, consider the case where $G = C_4$ is a circuit of length 4 and d satisfies: $d_{12} = d_{23} = 1$, $d_{34} = 3$ and $d_{14} = 0$. Then d cannot have an ℓ_2-realization since d has no metric completion $\hat{d}$, for otherwise one would have $\hat{d}_{13} \leq d_{12} + d_{23} = 2$ and $d_{34} \leq d_{14} + \hat{d}_{13}$ implies $\hat{d}_{13} \geq 3$, yielding a contradiction.

As a second example, consider the distance d defined on the graph $G = K_4$ by $d_{12} = d_{13} = d_{14} = 1$ and $d_{23} = d_{24} = d_{34} = 2$. Then d is a metric (and moreover d belongs to CUT_4, since all the facets of CUT_4 come from triangle inequalities). However d does not have an ℓ_2-realization; for otherwise one could find vectors $u_1 = 0$, u_2, u_3, u_4 satisfying $\|u_i\| = 1$ $(i = 2, 3, 4)$ and $\|u_i - u_j\| = 2$, i.e., $u_i^T u_j = -1$ $(i \neq j \in \{3, 4\})$, clearly a contradiction.

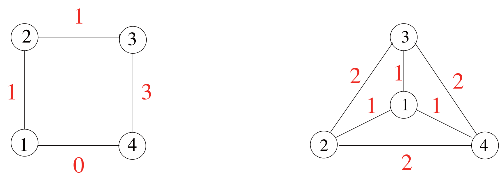

Fig. 2: Two distances with no ℓ_2-realization

The distance geometry problem is closely related to the *PSD matrix completion problem:* Decide whether a partially specified matrix can be completed to a positive semidefinite matrix. In other words, given $a \in \mathbb{R}^E$, is it possible to find a matrix $X \succeq 0$ such that $X_{ii} = 1$ $\forall i \in V$ and $X_{ij} = a_{ij}$ $\forall ij \in E$? That is, decide whether a belongs to the projection $\mathcal{E}(G)$ of the elliptope $\mathcal{E}_n$ onto the edge subspace $\mathbb{R}^E$. Again the exact complexity is not known. An obvious necessary condition for $a \in \mathcal{E}(G)$ is that the vector $\arccos(a) \in [0, \pi]^E$ can be completed to a spherical distance. More strongly, following our earlier discussion on 'angles of vectors', this spherical distance should satisfy all valid inequalities for the cut polytope. As an illustration, the following partial matrix:

$$\begin{pmatrix} 1 & \mathbf{1} & ? & -\mathbf{1} \\ \mathbf{1} & 1 & \mathbf{1} & ? \\ ? & \mathbf{1} & 1 & \mathbf{1} \\ -\mathbf{1} & ? & \mathbf{1} & 1 \end{pmatrix}$$

cannot be completed to a positive semidefinite matrix. This can be checked directly. Alternatively observe that assigning values $(0, 0, 0, \pi)$ $(= \arccos(1, 1, 1, -1))$ to the edges of C_4 gives a partial distance that cannot be completed to a spherical distance.

Grothendieck inequalities

Given a graph $G = (V = [n], E)$ and edge weights $w \in \mathbb{R}^E$, consider the quadratic integer problem:

$$\mathrm{IP}(w) := \max_{x_i \in \{\pm 1\}} \sum_{ij \in E} w_{ij} x_i x_j \tag{13}$$

and its semidefinite relaxation:

$$\mathrm{SDP}(w) := \max_{u_i \text{ unit vectors}} \sum_{ij \in E} w_{ij} u_i^T u_j. \tag{14}$$

Clearly, $\mathrm{IP}(w) \leq \mathrm{SDP}(w)$. Alon et al. [2006] showed the existence of a constant $K(G)$ such that

$$\mathrm{SDP}(w) \leq K(G)\ \mathrm{IP}(w) \quad \forall w \in \mathbb{R}^E.$$

Note that the program (13) corresponds to optimizing over the cut polytope $\mathrm{CUT}^\square(G)$, while (14) corresponds to optimizing over the elliptope $\mathcal{E}(G)$ (the respective projections of $\mathrm{CUT}_n^\square$ and $\mathcal{E}_n$ onto the edge subspace $\mathbb{R}^E$). Thus the constant $K(G)$ can be viewed as the integrality gap of the SDP relaxation or, in geometric terms, as the smallest dilation λ for which $\mathcal{E}(G) \subseteq \lambda\ \mathrm{CUT}(G)$. For instance, $K(G) = 3/2$ for $G = K_3$ (recall the picture of $\mathcal{E}(K_3)$ in Fig. 1).

The constant $K(G)$ is known as the *Grothendieck constant* of the graph G. Indeed, this concept goes back to work of Grothendieck[1] [1953] who showed that $K(G)$ is universally bounded for bipartite graphs. More precisely it is known that

$$1.6769 \leq \sup_{G \text{ bipartite}} K(G) < 1.8782.$$

These results were shown in the context of operator norms in functional analysis and can be used to design approximation algorithms, for instance, for the problem of computing the cut norm of a matrix.

More generally, Alon et al. [2006] show the existence of two absolute constants $C, C' > 0$ such that

$$C \log \omega(G) \leq K(G) \leq C' \log \vartheta(\overline{G}),$$

Alexander Grothendieck

[1] Grothendieck left the academic world since the 1980's and is now said to live in reclusion in southern France. In January 2010, he wrote a "Declaration d'intention de non-publication", where he asks that none of his work should be reproduced in whole or in part, and even further that libraries containing such copies of his work should remove them.

where $\omega(G)$ is the largest size of a clique in G and $\vartheta(\overline{G})$ is the celebrated *Lovász theta number*, known to be upper bounded by $\chi(G)$, the coloring number of G. As $\chi(G) = 2$ when G is a bipartite graph, this implies again Grothendieck's result, claiming that the Grothendieck constant is universally bounded for bipartite graphs.

Interest in the Grothendieck constant comes, in particular, from its relevance to approximation algorithms in complexity theory, as well as to non-local games in quantum information theory, briefly discussed below.

Bell inequalities in quantum information

A non-local game $\mathcal{G} = (\pi, V)$ involves three parties: the *referee* and two *players*, Alice and Bob. Data include a probability $\pi : S \times T \to [0,1]$ and a predicate $V : A \times B \times S \times T \to \{0,1\}$ giving the rules of the game. Here, A, B, S, T are given finite sets, S, T are the sets of 'questions' sent by the referee to Alice and Bob and A, B are the sets of 'answers' of Alice and Bob, respectively. The players and the referee know π and V. Before the start of the game Alice and Bob may decide on a strategy to play the game, but once the game has started they are not allowed to communicate.

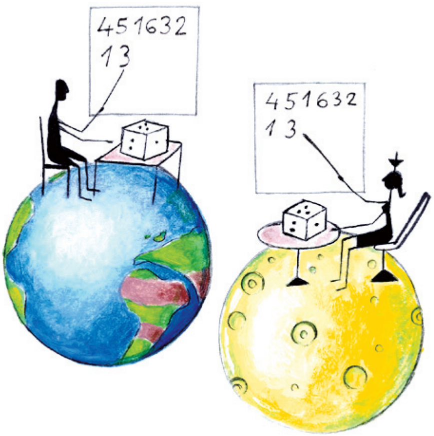

The game goes as follows:

1. The referee picks a pair of questions $(s,t) \in S \times T$ according to the ditribution π and sends question s to Alice and question t to Bob.

2. Alice selects her answer $a \in A$ and Bob selects his answer $b \in B$.

The players win the game if $V(a, b|s, t) = 1$ and loose otherwise. The aim of the players is to choose their own strategy so as to maximize their probability of winning the game. This probability can be computed with the following optimization problem:

$$\max_P \sum_{(s,t) \in S \times T} \pi(s,t) \sum_{(a,b) \in A \times B} V(a,b|s,t) P(a,b|s,t) \tag{15}$$

where P corresponds to all possible probability profiles, depending on the strategies adopted by the players.

In the *classical* deterministic case, the players simply choose a map $a : S \to A$ for Alice and $b : T \to B$ for Bob. The corresponding feasible probability profiles then form a polytope, whose bounding faces are supported by the so-called *Bell inequalities*. The optimal value of the corresponding optimization problem (15), called the *classical value* of the game, can be computed via LP.

In the non-classical setting, the players share an entangled strategy. It consists of a quantum system $\{A_s^a \mid a \in A, s \in S\}$ of bounded operators on a Hilbert space H_A for Alice, a quantum system $\{B_t^b \mid b \in B, t \in T\}$ of bounded operators on a Hilbert space H_B for Bob, and a shared entangled state Ψ which is a unit vector in $H_A \otimes H_B$. The operators A_s^a of Alice are required to be Hermitian projectors forming an orthogonal resolution of the identity, and analogously for the operators B_t^b of Bob; that is,

1. $A_s^a A_s^{a'} = \delta_{aa'} A_s^a$, $B_t^b B_t^{b'} = \delta_{bb'}$ $\forall a, a' \in A$, $\forall b, b' \in B$, $\forall s \in S$, $\forall t \in T$,
2. $\sum_{a \in A} A_s^a = I_A$, $\sum_{b \in B} B_t^b = I_B$ $\forall s \in S$, $\forall t \in T$.

Moreover, $P(a, b|s, t) = \langle \Psi, A_s^a \otimes B_b^t \ \Psi \rangle$. Hence the *quantum value* of the game can be computed via the optimization problem:

$$\max_{\Psi, A_s^a, B_t^b} \sum_{s,t} \pi(s,t) \sum_{a,b} V(a,b|s,t) \langle \Psi, A_s^a \otimes B_b^t \ \Psi \rangle, \tag{16}$$

where the maximum is taken over all possible unit vectors $\Psi \in H_A \otimes H_B$, A_s^a Hermitian operators on H_A, B_t^b Hermitian operators on H_B satisfying the above conditions 1-2.

In an XOR game, Alice and Bob have two possible answers, i.e., $A = B = \{0, 1\}$, and the predicate $V(a, b|s, t)$ depends only on the XOR value $a \oplus b$ of their answers. A classical result in quantum information theory is the theorem of Tsirelson [1987], showing that the optimization problem (16) can be reformulated as a semidefinite program, of the form

$$\max_{u_s, v_t \text{ unit vectors}} \sum_{s \in S, t \in T} \pi'(s,t) u_s^T v_t,$$

after suitably defining $\pi'(s, t)$. Note that this problem is of the form (14), where the graph is bipartite with bipartition (S, T). In a nutshell, for XOR games, the classical value of the game can be cast as an instance of the integer quadratic problem (13), while the quantum value of the game can be cast as the SDP relaxation (14), both for the case of bipartite graphs. Hence the Grothendieck's constant also permits to estimate the ratio between the classical and quantum values.

Moreover, the set of feasible probability profiles corresponding to classical strategies form a polytope, closely related to the cut polytope of the complete bipartite graph. In other words, Bell inequalities can be interpreted as valid inequalities for the cut polytope. The maximum violation of a Bell inequality that can be obtained when considering quantum strategies can be computed via semidefinite programming.

For general non-local games the program (16) can be reformulated as a polynomial optimization problem with *non-commutative variables* which can then be approximated by a hierarchy of semidefinite programs. This relies on a recent extension of the techniques developed by Lasserre and Parrilo around 2000 for classical (commutative) polynomial optimization problems, based on sums of squares of polynomials and the dual theory of moments.