Lecture 8 - Integer Linear Program: Integrality gaps and Erdős-Pósa property

Optimization and Approximation - ENS M1

Nicolas Bousquet

1 Integrality gap

Let (P) be an Integer Linear Program (ILP for short). The *fractional relaxation* of (P), denoted by (P*) is the LP where the Integrality constraint is replaced by a non-negativity constraint. In the particular case of $0-1$-Linear Program (ILP where variables can take either the value 0 or 1), the constraint $x \in \{0,1\}$ can be replaced by the two constraints $0 \leq x \leq 1$.

One may ask what is the difference between the optimal value of the ILP and of its fractional relaxation. The *integrality gap* of (P) is $\frac{OPT(P')}{OPT(P)}$ if P is a maximization function and $OPT(P)/OPT(P')$ if (P) is a minimization function. Note that the integrality gap is always at least one since an integral solution always is a solution of its fractional relaxation. The question then is: how large can be the integrality gap?

Bounding the integrality gap is in general an interesting property. Indeed, if the integrality gap is bounded, we often have a "natural" road-map to find an approximation algorithm:

- Find an optimal solution of the fractional relaxation (in polynomial time).
- Round the values of the variables of the fractional relaxation to transform it into a solution of the original LP.

Even if this method does not necessarily work each time, a LP with an arbitrarily integrality gap has less chance to have an approximation algorithm (at least using LP).

In what follows, we assume that P is a maximization function adn (D) a minimization function. *primal-dual gap* is the ration between the optimal value of the Integral Primal and between the integral Dual, i.e. $OPT(D)/OPT(P)$. Again if this ratio is bounded, a primal dual algorithm might provide a good approximation algorithm while otherwise such algorithms must not exist.

In what follows, we will illustrate integrality gaps and primal dual gaps on two classical problems: the minimum Hitting Set problem and the maximum Packing problem.

2 Hitting Set and Packing

Let $H = (V, E)$ be a hypergraph with no empty hyperedge. A *hitting set* of H, also called a *transversal*, is a subset X of vertices such that, for every hyperedge $e \in E$, at least one vertex of X is in e. In other words, a hitting set is a subset of vertices intersecting all the hyperedges. Note that the whole set of vertices V is a hitting set. The *transversality of* H, denoted by $\tau(H)$ is the minimum size of a hitting set of H. Determining if a hypergraph H has a hitting set of size at most k is an NP-hard problem, even for 2-uniform hypergraphs (one of the 21 Karp's NP-complete problems called Vertex Cover.
A subset of hyperedges $P \subseteq E$ is a *packing* of H if every vertex of H is in at most one hyperedge of P. In other words, all the hyperedges of P are vertex disjoint. The *packing number of* H, denoted by $\nu(H)$, is the maximum size of a packing of H. Note that a hypergraph H satisfies $\nu(H) = 1$ if all its hyperedges pairwise intersect. Determining if a hypergraph has a packing of size k is an NP-complete problem, even for 3-uniform hypergraphs (the problem is the 3-dimensional matching problem). Nevertheless for 2-uniform hypergraphs the problem is polynomial time solvable (indeed a packing of a 2-uniform

hypergraph is exactly a matching of the corresponding graph and a maximum matching can be found in polynomial time). When no confusion is possible, transversality and packing number will be denoted by τ and ν instead of $\tau(H)$ and $\nu(H)$.

Observation 1. *Every hypergraph with no empty hyperedge H satisfies*

$$\nu(H) \leq \tau(H).$$

Proof. Let P be a packing and X be a hitting set of H. Every hyperedge e of P satisfies $e \cap X \neq \varnothing$. And no vertex of X intersects more than one edge of the packing since $e \cap e' = \varnothing$ for every $e, e' \in P$. So $|X| \geq |P|$. The desired inequality is obtained by considering a maximum packing and a minimum hitting set. □

As we'll see later, it is a consequence of the duality of LP. In the following, we will consider the transversality and the packing number as linear programs. Let $H = (V, E)$ be a hypergraph.

> **Transversal (Integer) Linear Program:**
> **Variables:** A variable $x_v \in \mathbb{R}$ (or $\mathbb{N}$) for every $v \in V$.
> **Constraints:** For every hyperedge $e \in E$, $\sum_{v \in e} x_v \geq 1$.
> For every vertex v, $x_v \geq 0$.
> **Objective function:** minimize $\sum_{v \in V} x_v$.

Note that there are two types of constraints. The *hyperedge constraints* are the constraints on the hyperedges of H. The *non-negativity constraints* are the constraints on the vertices of H. There exists a bijection between the set of variables and the set of vertices of the hypergraph. By abuse of notation, and when no confusion is possible, we will say that the variables are the vertices of the hypergraph. The values of the variables can be seen as a weight function on the vertices of the hypergraph. A weight function $w : V \longrightarrow \mathbb{R}$ satisfies the transversal linear program if, when we give the value $w(v)$ to the variable x_v then all the constraints are satisfied. We will denote by $w(V)$ the sum of the weights of all the vertices.

Observation 2. *Let H be a hypergraph with no empty hyperedges. The optimal value of the Transversal Integer Linear Program of H equals $\tau(H)$.*

Proof. Let w be an optimal weight function $V \longrightarrow \mathbb{N}$ of the Transversal ILP. First note that each variable has either value zero or one. Indeed, assume by contradiction that some vertex v satisfies $w(v) \geq 2$. Consider the weight function w' where $w'(u) = w(u)$ for every $u \neq v$ and $w'(v) = w(v) - 1$. All the constraints are still satisfied. For the non-negativity constraints, it is immediate. For the hyperedge constraints, either $v \notin e$ and then we have $\sum_{u \in e} w'(u) = \sum_{u \in e} w(u) \geq 1$ since the function w is a solution of the transversal LP. Or $v \in e$, and then we have $\sum_{u \in e} w'(u) \geq w'(v) \geq 1$. So w' is also a solution of the transversal LP and $\sum_{u \in V} w'(u) < \sum_{u \in V} w(u)$, a contradiction with the optimality of w.

So w is a function $V \longrightarrow \{0, 1\}$. Let X be the subset of vertices v such that $w(v) = 1$. Since $\sum_{u \in e} w(u) \geq 1$ for every hyperedge e, we have $X \cap e \neq \varnothing$, *i.e.*$X$ is a hitting set. Conversely, every hitting set can be transformed into a solution of the Transversal ILP by giving weight one to the vertices of the hitting set and zero to the others. □

Note that if the hypergraph contains an empty hyperedge, then the Transversal LP cannot be satisfied since there is a constraint $0 \geq 1$. Indeed, no vertex of the hypergraph can intersect the empty hyperedge. Let us now define the packing LP of a hypergraph $H = (V, E)$.

> **Packing (Integer) Linear Program:**
> **Variables:** A variable $x_e \in \mathbb{R}$ (or $\mathbb{N}$) for each hyperedge e.
> **Constraints:** For every vertex v, $\sum_{e | v \in e} x_e \leq 1$.
> For every hyperedge e, $x_e \geq 0$.
> **Objective function:** maximize $\sum_{x_e \in E} x_e$.

Note that if the hypergraph H contains an empty hyperedge then the optimal value of the Packing Linear Program of H is infinite. Indeed, we can give to the empty hyperedge an arbitrarily large value since no constraint uses it.

Observation 3. *Let H be a hypergraph with no empty hyperedge. The optimal value of the Packing Integer Linear Program of H equals $\nu(H)$.*

Proof. First note that every variable has value 0 or 1. Indeed if a variable x_e has value at least two, then every vertex $v \in e$ would satisfy $\sum_{e'|v \in e'} x'_e \geq 2$, a contradiction. Such a vertex v exists since H does not contain any empty hyperedge.
So the Packing ILP can be seen as a weight function $w : E \longrightarrow \{0, 1\}$. Let us denote by P the set of hyperedges e such that $w(e) = 1$. Since every vertex v satisfies $\sum_{e|v \in e} w(e) \leq 1$, every vertex is in at most one hyperedge of P. In other words the hyperedges of P are vertex disjoint, *i.e.*P is a packing. □

If the hypergraph contains an empty hyperedge, then we can put it in the packing (it does not intersect any other hyperedge, so all the optimal packings contain it) and compute the packing of the remaining hypergraph via linear programming: containing an empty hyperedge does not avoid us to apply linear programming techniques.

Theorem 4. *Transversal Linear Program and Packing Linear Program are dual.*

Sketch of proof. Let H be a hypergraph. Let us denote by A the matrix constraint, by b the constraint vector and by c the objective function of the Transversal LP. So the Transversal LP can be written as $\min {}^t cx$ under the constraint $Ax \leq b$.

In the dual linear program, there is a variable associated to each constraint of the primal linear program. In other words, there is a variable associated to every hyperedge constraint (we will denote it by x_e) and to every positivity constraint (we will denote it by y_v). In the following, we denote by y the concatenation of the vectors x_e and y_v. Let us first describe the objective function of the dual LP. The objective function of the dual LP consists in maximizing ${}^t by$ (since the dual of a minimization problem is a maximization problem). The vector b equals one on the hyperedge constraints (since every hyperedge must have size at least one) and equals zero on the positivity constraints (since every variable must have non negative weight). So the objective function of the dual LP is $\max(\sum_{e \in E} x_e)$.

The vector c is a vector of ones (since we sum the values of all the variables in the objective function of the primal LP). The coefficients of the matrix A are 0 and 1. On the lines corresponding to hyperedge constraints, the coefficient $A(j, i) = 1$ if the vertex i is in the hyperedge e_j. On the lines corresponding to positivity constraints, all the coefficients equal zero except in the j-th column for the positivity constraint of the vertex v_j. The constraints of the dual LP correspond to the columns (*i.e.*to the vertices) of the primal LP. The dual constraint corresponding to the vertex v is $(\sum_{e|v \in e} x_e) + y_v = 1$. The dual linear program has the following variables and constraints.

Variables: A variable x_e for each hyperedge e,
A variable y_v for every vertex v.
Constraints: For every vertex v, $y_v + \sum_{e|v \in e} x_e = 1$.
For every hyperedge e, $x_e \geq 0$.
For every vertex v, $y_v \geq 0$.

Note that each variable y_v appears in exactly one constraint (the constraint of the vertex v) and does not appear in the objective function. So we can forget these variables and replace the equalities by inequalities. In other words, the vertex constraint for the vertex v becomes $\sum_{e|v \in e} x_e \leq 1$. It is equivalent to the original one since if we add y_v and give it the value $1 - \sum_{e|v \in e} x_e$, the original constraint is satisfied. The resulting LP is exactly the Packing LP. □

In the following we will denote by τ^* and ν^* the fractional relaxation of respectively τ and ν. In other words, τ^* (resp. ν^*) is the optimal value of the Transversal (resp. Packing) Linear Program in real numbers. The duality theorems ensures that

Theorem 5. *Every hypergraph H with no empty hyperedge satisfies :*

$$\nu \leq \nu^* = \tau^* \leq \tau.$$

Before illustrating the integrality gap of transversal and packing LP, let us first provide a simple but useful observation.

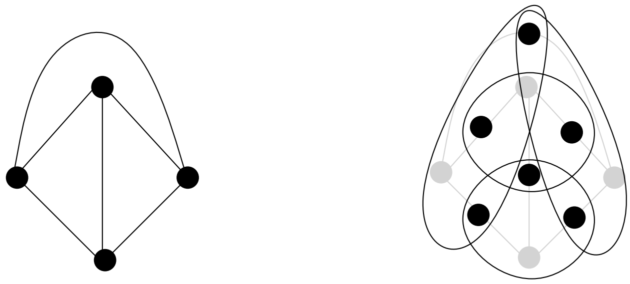

Figure 1: Illustration of the proof of Lemma 8. At the left a K_4. At the right, the hypergraph constructed from the edges of K_4.

Observation 6. *Let H be a hypergraph with no empty hyperedge and let c be a positive constant. If every hyperedge contains at least $c \cdot n$ vertices, then $\tau^* \leq 1/c$.*

Proof. Let w be the weight function $V \longrightarrow \mathbb{R}$ such that $w(v) = 1/(cn)$ for every $v \in V$. Every hyperedge has weight at least 1 (since every hyperedge contains at least cn vertices and each vertex has weight $1/cn$), so all the constraints are satisfied. And the total weight of the vertex set is $n/cn = 1/c$. □

Lemma 7. *The gap between τ and τ^* can be arbitrarily large.*

Proof. Let $\mathcal{U}_{n,2n}$ be the complete n-uniform hypergraph on $2n$ vertices (*i.e.* where all the subsets of size n are hyperedges). Since every hyperedge contains half of the vertices, Observation 6 ensures that $\tau^* \leq 2$. On the contrary, we have $\tau \geq n + 1$. Otherwise the complement of a hitting set would have size at least n, and then would contain a hyperedge (since every subset of size n is a hyperedge), a contradiction. □

Lemma 7 refines the fact that the gap between τ and ν can be arbitrarily large by Theorem 5.

Lemma 8. *The gap between ν and ν^* can be arbitrarily large.*

Proof. Let K_n be a clique on n vertices. Construct the following hypergraph H_n. The vertices of H_n are the edges of K_n. For every vertex v, create the hyperedge e_v containing all the edges adjacent to v. The hypergraph H_4 is represented on Figure 1.

Consider the following weight function w on the hyperedges which associates $1/2$ to every hyperedge of the hypergraph. The total weight is $n/2$ since there are n hyperedges in the hypergraph H_n and each hyperedge has weight $1/2$. The constraints are satisfied since every vertex of H_n is in two hyperedges (the edge uv of K_n is only in the hyperedges e_u and e_v). So $\nu^*(H_n) \geq n/2$. On the contrary, we have $\nu(H_n) = 1$. Indeed, for every pair of vertices u, v, the hyperedges e_u and e_v intersect on uv (the edge uv exists since the original graph is a clique). □

3 Erdős-Pósa property

The hitting set problem and the maximum packing problem permits to illustrate that:

- The integrality gaps can be unbounded.
- The primal-dual gap can be unbounded.

A class $\mathcal{P}$ of problems has the *Erdős-Pósa property* if there exists a function f such that for any instance of $\mathcal{P}$, we opt $OPT(D) \leq f(OPT(P))$. Since the integrality gap can be unbounded in general (e.g. for the hitting set problem), there exist problems which DO NOT have the Erdős-Pósa property.

Example: Vertex Cover

Exercise 1.

- *Show that the Vertex-Cover problesms and the Maximum Matching problems are dual.*
- *Show the the primal-dual gap is at most two.*

Generally speaking, if the sizes of the hyperedges can be bounded by a constant, then the class has the Erdős-Pósa property.

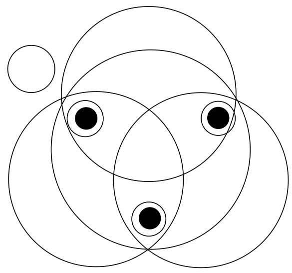

Figure 2: A shattered set of size 3.

Erdős-Posá property and integrality gap. Note that if the class of hypergraphs have the Erdos-Posa property, then the integrality gap is indeed bounded by the same function. Indeed, the weak duality theorem ensures that: $\nu \leq nu^* = \tau^* \leq \tau$. So if $\tau \leq f(\nu)$ then $\tau \leq f(\tau^*)$ for the same function f. However, if the integrality gap is bounded (e.g. if $\tau \leq 2 \cdot \tau^*$) it does not imply that the other integrality gap, and then the gap between τ and ν is bounded.

Exercise 2. *Search in the lecture an example where an integrality gap is bounded and the primal-dual gap is unbounded.*

4 Shattered sets

Let H be a hypergraph. A subset X of vertices is *shattered* if for every subset $X' \subseteq X$, there exists a hyperedge e such that $e \cap X = X'$. Note that the empty set is always shattered. Figure 2 represents a shattered set of size 3. In other words, a shattered set is a subset X of vertices such that all the possible traces of H on X exist. Yet differently, the hyperedges intersect in all the possible ways the set X. In the following, we denote by $sh(H)$ the number of shattered sets in H. Since a shattered set is a subset of vertices, we have $sh(H) \leq 2^n$.
A shattered set is a witness of the local complexity of the hypergraph: a hypergraph is complex on a set X if all the traces on X exist. The interest of shattered sets (and of VC-dimension) is that a bounded local complexity provides several general properties on the hypergraph, for instance on the number of hyperedges (Theorem 12) or on the size of hitting sets.

The *Vapnik-Chervonenkis dimension*, or *VC-dimension* for short, is the maximum size of a shattered set. In the following we will only deal with hypergraphs of VC-dimension at least one. Indeed as long as a hypergraph contains two (distinct) hyperedges, it contains a shattered set of size one. A hypergraph has VC-dimension at least d if and only if it contains $\mathcal{C}_d$ as a subhypergraph. Let us denote by $vc(H)$ the VC-dimension of H. Given a family of hypergraphs $\mathcal{H}$, the VC-dimension of $\mathcal{H}$ is the largest VC-dimension of a hypergraph of $\mathcal{H}$. The VC-dimension was first introduced by Vapnik and Chervonenkis in 1971.

4.1 Opening properties

Observation 9. *Every hypergraph H with m hyperedges satisfies $VC(H) \leq \lfloor \log(m) \rfloor$.*

Observation 10. *Every hypergraph H of VC-dimension d satisfies $sh(H) \leq \sum_{i=0}^{d} \binom{n}{i}$.*

Proof. The proof is straightforward: if the VC-dimension is at most d, then only sets of size at most d can be shattered. And the number of such sets is $\sum_{i=0}^{d} \binom{n}{i}$. □

Note that the number of shattered sets gives more precise information than the VC-dimension. Indeed, Observation 10 ensures that as long as the VC-dimension is at most d, the number of shattered sets is at most $\sum_{i=0}^{d} \binom{n}{i}$. Nevertheless, a reverse function which does not depend on n does not exists.

Observation 11. *For every n, there exist hypergraphs with n vertices and at most n hyperedges and at most n shattered sets with VC-dimension $\lfloor \log n \rfloor$.*

Proof. Let V be a set of size n and X be a subset of V of size $\lfloor \log n \rfloor$. Denote by H the hypergraph with vertex set V whose hyperedges are all possible subsets of X. By construction, the shattered sets are the subsets of X. Indeed, since no hyperedge contains a vertex which is not in X, no set containing a vertex which is not in X is shattered. Finally, $vc(H) = \lfloor \log n \rfloor$ (since X is shattered) but $sh(H) = 2^{\lfloor \log n \rfloor} \leq n$. □

Note that the VC-dimension of a hypergraph $H = (V, E)$ is at most $\lfloor \ln |E| \rfloor$. Indeed, consider a shattered set X. All the possible traces exist on X. Since there are $2^{|X|}$ traces on every set X (all the possible subsets of X), the hypergraph contains at least $2^{|X|}$ hyperedges.

4.2 First examples

Geometrical hypergraphs Intersections of geometrical objects are often "simple", in the sense that they are constrained because of the topology and geometry of the space. Since the VC-dimension catches the complexity of the intersections of objects, it is natural to think that the VC-dimension must be bounded above in many cases.

Consider one of the easiest class of intersection of objects: the intersection of intervals in the real line (*i.e.*in $\mathbb{R}^1$). Vertices are points of the real line and intervals are represented by hyperedges. Let x_1, x_2, x_3 be three real numbers such that $x_1 \leq x_2 \leq x_3$. Every interval containing both x_1 and x_3 must contain x_2 (since an interval is convex). So $\{x_1, x_2, x_3\}$ cannot be shattered since no interval e satisfies $\{x_1, x_2, x_3\} \cap e = \{x_1, x_3\}$. Then the VC-dimension is at most 2.

More generally, the VC-dimension of intersection of axis-parallel d-dimensional rectangles is at most $2d$. A d-dimensional rectangle hypergraph is a hypergraph where the set of vertices V is a set of points of $\mathbb{R}^d$ and a hyperedge corresponds to the intersection of an axis-parallel d-dimensional rectangle with V. Let X be a set of size $2d + 1$. Consider the set S containing all the vertices of X which are maximum or minimum for at least one of the d coordinates. If there are several minimum or maximum values, choose arbitrarily one of them. The set S contains at most $2d$ points (note that a same point can be maximum or minimum for several coordinates but it does not matter). Any rectangle containing all the points of S contains all the points of X which are between the first and the last point of each coordinate. In other words, any rectangle containing all the vertices of S also contains the whole set X. Hence X cannot be shattered since no rectangle has trace S.

5 VC-dimension and number of hyperedges

This Section is devoted to proving that any hypergraph of bounded VC-dimension has a polynomial number of hyperedges. This result is known as Sauer's Lemma. Recall that general hypergraphs can have up to 2^n hyperedges. In order to prove it, we prove a stronger statement: the number of hyperedges can be bounded above by the number of shattered sets. We propose a classical proof of this result using a shifting argument.

5.1 Sauer's Lemma

A hypergraph is simple if it does not contain twice the same hyperedge.

Lemma 12 (Sauer's Lemma). *Let $H = (V, E)$ be a simple hypergraph of VC-dimension d. For every set $X \subseteq V$, the number of (distinct) traces of E on X is at most $\sum_{i=0}^{d} \binom{|X|}{i}$. In particular, we have*

$$|E| \leq \sum_{i=0}^{d} \binom{n}{i}$$

Sketch of the proof. The proof is a proof by induction on $d + n$ where the base cases are the cases $d = 0$ and $n = 0$. If $d = 0$ then the number of hyperedges is at most 0. Indeed if H contains two hyperedges, since the hypergraph is simple, they are distinct. Thus there exists a vertex v which is in the first and not in the second. Thus, the set $\{v\}$ is shattered, a contradiction. So if $d = 0$ then $|E| \leq 1$ and the conclusion holds. If $n = 0$ then there is at most one hyperedge which is the empty set.

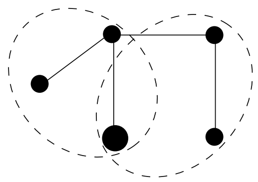

Figure 3: The big vertex is a $1/2$-net (with the uniform measure) but not a hitting set

Assume now that $d \geq 1$ and $n \geq 1$. Let v be a vertex and let E_1, E_2 be a partition of E defined as follows:

$$E_1 = \{e \text{ such that } v \in e \text{ and } e \setminus v \in E\}$$
$$E_2 = \{e \text{ such that } \exists e' \in E \setminus E_1, e = e' \setminus v\}$$

Let us define H_1 (resp. H_2) as the hypergraph on vertex set $V \setminus v$ and with edge set E_1 (resp. E_2). Let us prove that (i) both hypergraphs are simple, (ii) the VC-dimension of H_1 is at most $d-1$.

(i) By definition of E_1, the hypergraph H_1 is simple. Assume by contradiction that H_2 is not simple. So there exists two hyperedges e_1 and e_2 of H that gives the same hyperedge e' in H_2. Since $e' = e_1 \setminus v$ and $e' = e_2 \setminus v$ and $e_1 \neq e_2$ (since H is simple), we have $e_1 = e_2 \cup v$ (or the opposite). But then $e_1 \in E_1$, a contradiction.

(ii) Let us now prove that H_1 has VC-dimension at most $d-1$. In order to prove it, you just have to show that if X is shattered in H_1 then $X \cup v$ is shattered in H.
Exercise: Prove it !

Now we apply the induction hypothesis on both sides, and we obtain:

$$|E| = |E_1| + |E_2| \leq \sum_{i=0}^{d-1} \binom{n-1}{i} + \sum_{i=0}^{d} \binom{n-1}{i} \leq \sum_{i=0}^{d} \binom{n}{i}$$

The first inequality holds by induction and the second is a basic calculation of binomial functions. □

Note that, in many cases, we do not need the exact upper bound on the number of hyperedges, but just an upper bound. In these cases, we will implicitly use the following inequality:

$$\sum_{i=0}^{d} \binom{|X|}{i} \leq k^{d+1}.$$

Note that, since $sh(H[E \setminus e]) \leq sh(H)$ (a shattered set is still shattered in a larger hypergraph), we have $sh(H[E \setminus e])$ equals $|E \setminus e|$ or equals $|E \setminus e| + 1$ for every edge of a s-extremal hypergraph. Mészáros and Rónyai asked if there always exist an edge satisfying the equality.

In the following, we study a little bit later the structure of hypergraphs such that $|E(H)| = \sum_{i=0}^{d} \binom{n}{i}$.

6 VC-dimension and integrality gap

6.1 Upper bounds on the size of ϵ-nets

A *measure* of a hypergraph is a weight (*i.e.*non negative) function on the vertex set such that the sum of the weights equals one. Let H be a hypergraph and μ be a measure on the vertex set of H. An *ϵ-net* is a subset of vertices X such that every hyperedge of weight at least ϵ intersects X. The *uniform measure* is the measure where all the vertices are given the same weight (*i.e.*weight $1/|V|$). In this case, an ϵ-net is a subset of vertices intersecting every hyperedge of size at least ϵn. In Figure 3, the big vertex is an $1/2$-net (for the uniform measure) since it intersects all the hyperedges of size at least $1/2$. Note that it is not a hitting set. On the contrary, a $2/5$-net would be a hitting set of the hypergraph. More generally, let c be the minimum size of a hyperedge, a c/n-net is a hitting set (for the uniform measure). A key result of Haussler and Welzl in VC-dimension theory bounds above the minimum size of an ϵ-net in function of the VC-dimension and of ϵ.

Theorem 13. *[Haussler, Welzl] Every hypergraph of VC-dimension d has an ϵ-net of size $\mathcal{O}(\frac{d\ln(d/\epsilon)}{\epsilon})$.*

Sketch of the proof. The proof of Theorem 13 is tricky. In the following the main steps of the proof are presented, but for simplicity, technical details and calculations are omitted. For simplicity, we assume that the measure is uniform. We can assume without loss of generality that all the hyperedges have size at least ϵn. Indeed the others do not have to intersect the ϵ-net so we omit them.

Draw a random vertex subset X of size $s := C \cdot (d/\epsilon)\ln(d/\epsilon)$ where C is a positive constant not detailed here. The proof is devoted to showing that X is an ϵ-net with positive probability. For every hyperedge e and every vertex $x \in X$, we have $\mathbb{P}(x \in e) \geq \epsilon$ since e contains at least ϵn vertices and X is drawn randomly. Since vertices of X are picked independently, the average number of vertices of X in e is at least $s \cdot \epsilon$. So, for every $e \in E$, Tchebychev inequality ensures that $\mathbb{P}(|e \cap X| \geq s\epsilon/2) \geq 1/2$. A hyperedge e *heavily intersects* X if $|e \cap X| \geq s\epsilon/2$.

Let us call E_0 the event "there exists a hyperedge which is not intersected by X". Note that E_0 is the complement event of our objective. Since we want to show that the complement of E_0 has a strictly positive probability, let us prove that $\mathbb{P}(E_0) < 1$. Draw randomly a second vertex subset Y of size s. Call E_1 the event "there exists a hyperedge which is not intersected by X and which heavily intersects Y". We have $\mathbb{P}(E_1) \leq \mathbb{P}(E_0)$ since E_1 requires E_0. And we also have $\mathbb{P}(E_1) \geq 1/2\mathbb{P}(E_0)$. Indeed, informally, if X is not an ϵ-net, then a hyperedge e is not intersected by X and e has probability at least $1/2$ to be heavily intersected by Y (since X and Y are drawn independently).

The rest part of the proof consists in proving that $\mathbb{P}(E_1) < 1/2$, which implies that $\mathbb{P}(E_0) < 1$. It finally ensures that there is a positive probability that X is an ϵ-net, *i.e.*at least one set of size s is an ϵ-net.
Let $k = s\epsilon/2$. Let A be a set of size $2s$. Keep in mind that the first half of A corresponds to X and the second half corresponds to B. Let us now draw a subset Z of size k in A. We want to compute the probability P that no vertex of Z appears in the first part of A (*i.e.*the probability that a hyperedge does not intersect X but heavily intersects Y). It is equivalent to force a set Z of size k to be in the second part of an equal bipartition of A. There are $\binom{2s-k}{s}$ bipartitions in which Z is in the second part. And the total number of bipartitions equals $\binom{2s}{s}$. So:

$$P = \frac{\binom{2s-k}{s}}{\binom{2s}{s}} \leq (1/\epsilon)^{-Cd/4}.$$

The last inequality is not immediate but can be obtained by a non trivial sequence of calculations (no deep argument is used at this point). Note that, until now, no argument based on VC-dimension has been used. Let us now link the events E_0 and E_1. Without the VC-dimension, the unique thing we can claim is that, since there are m hyperedges, $\mathbb{P}(E_1) \leq (1/\epsilon)^{-Cd/4}m$. But since the VC-dimension is at most d, Lemma 12 ensures that the number of traces on A is at most $|A|^{d+1}$. In other words, we have $\mathbb{P}(E_1) \leq (1/\epsilon)^{-Cd/4} \cdot (2s)^{d+1}$. A last sequence of non-trivial calculations ensures that $\mathbb{P}(E_1) < 1/2$. So $\mathbb{P}(E_0) < 1$, which achieves the proof. □

The proof ensures that when we pick randomly vertices, then the probability of finding an ϵ-net is positive. Hence it provides an randomized approximation algorithm for finding transversal in a hypergraph of bounded VC-dimension. Deterministic proofs (and algorithms) of Theorem 13 also exist (a derandomized proof was for instance proposed by Matoušek).

Note that the size of an ϵ-net can be at least $1/\epsilon$. Indeed, the size of an ϵ-net of a hypergraph containing $1/\epsilon$ disjoint hyperedges of size ϵn is at least $1/\epsilon$ (an ϵ-net must intersect each of the $1/\epsilon$ disjoint hyperedges).

Theorem 14. *Let $d \geq 2$. Denote by $s(d,n)$ the maximum size of a minimum ϵ-net over the set of hypergraphs of VC-dimension d on n vertices. We have:*

$$\frac{(d - 2 + \frac{1}{d+2} + o(1))\ln(1/\epsilon)}{\epsilon} \leq s(d,n) \leq \frac{(d + o(1))\ln(1/\epsilon)}{\epsilon}.$$

where $o(1)$ is a function which tends to zero when n tends to infinity.

Theorem 13 can be rephrased in order to obtain a bounded gap between τ and τ^*. In the following we will exclusively use this formulation instead of the (altough stronger) formulation with ϵ-nets.

Corollary 15. *Every hypergraph of VC-dimension d satisfies:*

$$\tau \leq \mathcal{O}(d\tau^* \ln(d\tau^*)).$$

Proof. Let w be a weight function on the vertex set corresponding to an optimal solution of the transversal linear program, *i.e.*$\sum_{x \in V} w(x) = \tau^*$ and for every hyperedge e, $\sum_{v \in e} w(e) \geq 1$. Let μ be the weight function such that each vertex x satisfies $\mu(x) = w(x)/\tau^*$. Note that $\sum_{x \in V} \mu(x) = 1$, *i.e.*$\mu$ is a measure. Every hyperedge has weight at least $1/\tau^*$ for μ since every hyperedge has weight at least one for w. So a $(1/\tau^*)$-net is a subset of vertices which intersects every hyperedge, *i.e.*a hitting set. Theorem 13 ensures that there is a hitting set of size $\mathcal{O}(d\tau^* \ln(d\tau^*))$. □

Corollary 15 ensures that the integrality gap between τ and τ^* is bounded if the VC-dimension is bounded. Previous statements do not explicit constants. The following theorem explicit them.

Theorem 16. *Every hypergraph H with $|E(H)| \geq 2$ of VC-dimension d satisfies*

$$\tau \leq 2d\tau^* \log(11\tau^*).$$

Erdős-Pósa property Since the integrality gap between τ and τ^* is bounded, Theorem 16 raises the following question: does the same hold for ν and ν^*? Recall that ν denotes the packing number and ν^* its fractional relaxation. It would imply immediately the Erdős-Pósa property since $\tau^* = \nu^*$ by Theorem 5. Unfortunately, the following example ensures that the integrality gap for ν is not bounded, even for hypergraphs of VC-dimension 2.

Lemma 17. *There exist hypergraphs of VC-dimension 2 such that $\nu = 1$ and $\nu^* = \mathcal{O}(\sqrt{n})$.*

Proof. The counter-example is the same as in Lemma 8 which provides an arbitrarily large gap between ν and ν^*. Let K_n be a clique on n vertices. Construct a hypergraph H_n such that vertices of H_n correspond to edges of K_n. Hyperedges of H_n correspond to the subsets of edges adjacent to a same vertex. We have seen in Lemma 8 that $\nu(H_n) = 1$ and $\nu^*(H_n) = \mathcal{O}(n)$. In addition, every vertex of H_n is contained in exactly two hyperedges, so no set of size 3 is shattered. Indeed, given a set X of size k, exactly 2^{k-1} subsets of X contain a fixed element x. So if a set of size 3 is shattered, at least 4 hyperedges contain each of the shattered vertices, a contradiction. □

7 2VC dimension and Erdős-Posá property

A subset of vertices X of a hypergraph is 2-*shattered*, if for every subset $X' \subseteq X$ of size 2, there exists a hyperedge e such that $e \cap X = X'$. The 2*VC-dimension* is the maximum size of a 2-shattered set. The *dual* 2*VC-dimension* is the maximum size of a 2-shattered set in the dual hypergraph. Note that dual 2VC-dimension is equivalent to the existence of a 2-complete Venn diagram, where 2-complete means that a vertex is contained in all the possible intersections of size 2. In other words a set $e_1, \ldots, e_d$ forms a 2-complete Venn diagram if for every i, j there exists $x_{i,j}$ which is in $e_i \cap e_j$ and which is in no other e_k with $k \neq i, j$.

First note that the VC-dimension is not larger than the 2VC-dimension. Indeed, a shattered set is in particular a 2-shattered set. As a consequence, Lemma 12 ensures that the maximum number of hyperedges of a hypergraph of 2VC-dimension d is at most $\sum_{i=0}^{d} \binom{n}{i}$. One can naturally ask for the existence of a function linking VC-dimension and 2VC-dimension. The following observation states that such a function does not exist.

We denote by $\mathcal{U}_{2,n}$ is the complete 2-uniform hypergraph, ie the hypergraph with all the hyperedges of size 2 and no hyperedges of larger size (it can be interpreted as a "classical" clique).

Observation 18. *For every $n \geq 4$, $\mathcal{U}_{2,n}$ has VC-dimension 2 and 2VC-dimension n.*

Proof. The VC-dimension of $\mathcal{U}_{2,n}$ equals 2 (as $\max(2, n-2) = 2$ since $n \geq 4$). And the 2VC-dimension of $\mathcal{U}_{2,n}$ is n since for every pair of vertices, there exists a hyperedge containing both of them and containing no other vertices of the graph. □

The following theorem, due to Ding, Seymour and Winkler is the main result of this Section.

Theorem 19 (Ding, Seymour, Winkler). *Every hypergraph H of dual $2VC$-dimension d satisfies:*

$$\tau \leq 11d^2(d+\nu+3)\binom{\nu+d}{\nu}^2$$