Discrete Mathematics and Theoretical Computer Science DMTCS vol. **23**:2, 2022, #5

On the Boolean dimension of a graph and other related parameters

Maurice Pouzet[1,2] Hamza Si Kaddour[1] Bhalchandra D. Thatte[3*]

[1] *Univ. Lyon, Université Claude-Bernard Lyon1, CNRS UMR 5208, Institut Camille Jordan, 43, Bd. du 11 Novembre 1918, 69622 Villeurbanne, France*
[2] *Department of Mathematics and Statistics, University of Calgary, Calgary, Alberta, Canada*
[3] *Departamento de Matematica, Universidade Federal de Minas Gerais (UFMG), Av. Antonio Carlos, 6627, Caixa Postal 702, Região Pampulha, Belo Horizonte - MG, CEP: 31270-901, Brasil*

received 3rd May 2021, revised 4th Apr. 2022, accepted 2nd Aug. 2022.

We present the Boolean dimension of a graph, we relate it with the notions of inner, geometric and symplectic dimensions, and with the rank and minrank of a graph. We obtain an exact formula for the Boolean dimension of a tree in terms of a certain star decomposition. We relate the Boolean dimension with the inversion index of a tournament.

Keywords: graphs, Boolean sum, symplectic dimension, geometric dimension, tournaments, inversion index

1 Presentation and preliminaries

We define the notion of Boolean dimension of a graph, as it appears in Belkhechine et al. (2010) (see also (Belkhechine, 2009; Belkhechine et al., 2012)). We present the notions of geometric and symplectic dimensions, and the rank and minrank of a graph, which have been considered earlier. When finite, the Boolean dimension corresponds to the inner dimension; it plays an intermediate role between the geometric and symplectic dimensions, and does not seem to have been considered earlier. The notion of Boolean dimension was introduced in order to study tournaments and their reduction to acyclic tournaments by means of inversions. The key concept is the inversion index of a tournament (Belkhechine, 2009; Belkhechine et al., 2010, 2012) presented in Section 3. Our main results are an exact formula for the Boolean dimension of a tree in terms of a certain star decomposition (Theorem 2.9) and the computation of the inversion index of an acyclic sum of 3-cycles (Theorem 3.7).

Notations in this paper are quite elementary. The *diagonal* of a set X is the set $\Delta_X := \{(x,x) : x \in X\}$. We denote by $\wp(X)$ the collection of subsets of X, by X^m the set of m-tuples $(x_1, \dots, x_m)$ of elements in X, by $[X]^m$ the m-element subsets of X, and by $[X]^{<\omega}$ the collection of finite subsets of X. The

*Supported by CAPES Brazil (Processo: 88887.364676/2019-00); the stay of this author was supported by LABEX MILYON (ANR-10-LABX-0070) of Université de Lyon within the program "Investissements d'Avenir (ANR-11-IDEX-0007)" operated by the French National Research Agency (ANR)

ISSN 1365–8050

cardinality of X is denoted by $|X|$. We denote by $\aleph_0$ the first infinite cardinal, by $\aleph_1$ the first uncountable cardinal, and by ω_1 the first uncountable ordinal. A cardinal κ is *regular* if no set X of cardinal κ can be divided in strictly less than κ subsets, all of cardinality strictly less than κ. If κ denotes a cardinal, 2^κ is the cardinality of the power set $\wp(X)$ of any set X of cardinality κ. If κ is an infinite cardinal, we set $\log_2(\kappa)$ for the least cardinal μ such that $\kappa \leq 2^\mu$. We note that for an uncountable cardinal κ the equality $\log_2(2^\kappa) = \kappa$ may require some set theoretical axioms, such as the Generalized Continuum Hypothesis (GCH). If κ is an integer, we use $\log_2(\kappa)$ in the ordinary sense, hence the least integer μ such that $\kappa \leq 2^\mu$ is $\lceil \log_2 \kappa \rceil$. We refer the reader to Jech (2003) and Kunen (2011) for further background about axioms of set theory if needed.

The graphs we consider are undirected and have no loops. They do not need to be finite, but our main results are for finite graphs. A *graph* is a pair (V, E) where E is a subset of $[V]^2$, the set of 2-element subsets of V. Elements of V are the *vertices* and elements of E are the *edges*. Given a graph G, we denote by $V(G)$ its vertex set and by $E(G)$ its edge set. For $u, v \in V(G)$, we write $u \sim v$ and say that u and v are *adjacent* if there is an edge joining u and v. The *neighbourhood* of a vertex u in G is the set $N_G(u)$ of vertices adjacent to u. The *degree* $d_G(u)$ of a vertex u is the cardinality of $N_G(u)$. If X is a subset of $V(G)$, the *subgraph of G induced by X* is $G_{\restriction X} := (X, E \cap [X]^2)$. A *clique* in a graph G is a set X of vertices such that any two distinct vertices in X are adjacent. If X is a subset of a set V, we set $K_X^V := (V, [X]^2)$; we say also that this graph is a clique.

1.1 The Boolean sum of graphs and the Boolean dimension of a graph

Let $(G_i)_{i\in I}$ be a family of graphs, all with the same vertex set V. The *Boolean sum* of this family is the graph, denoted by $\dot{+}(G_i)_{i\in I}$, with vertex set V such that an unordered pair $e := \{x, y\}$ of distinct elements of V is an edge if and only if it belongs to a finite and odd number of $E(G_i)$. If the family consists of two elements, say $(G_i)_{i\in\{0,1\}}$ we denote this sum by $G_0 \dot{+} G_1$. This is an associative operation (but, beware, infinite sums are not associative). If each $E(G_i)$ is the set of edges of some clique C_i, we say (a bit improperly) that $\dot{+}(G_i)_{i\in I}$ is a sum of cliques. We define the *Boolean dimension* of a graph G, which we denote by $\dim_{\mathrm{Bool}}(G)$, as the least cardinal κ such that G is a Boolean sum of κ cliques. In all, $\dim_{\mathrm{Bool}}(G) = \kappa$ if there is a family of κ subsets $(C_i)_{i\in I}$ of $V(G)$, and not less, such that an unordered pair $e := \{x, y\}$ of distinct elements is an edge of G if and only if it is included in a finite and odd number of C_i's.

A *Boolean representation* of a graph G in a set E is a map $f : V(G) \to \wp(E)$ such that an unordered pair $e := \{x, y\}$ of distinct elements is an edge of G if and only if the intersection $f(x) \cap f(y)$ is finite and has an odd number of elements.

Example 1.1. Let G be a graph. For a vertex $x \in V(G)$, let $E_G(x) := \{e \in E(G) : x \in e\}$. Set $E := E(G)$. Then the map $f : V(G) \to \wp(E)$ defined by $f(x) := E_G(x)$ is a Boolean representation. Indeed, for every 2-element subset $e := \{x, y\}$ of $V(G)$, the intersection $f(x) \cap f(y)$ has one element if and only if $e \in E(G)$, otherwise it is empty.

The following result is immediate, still it has some importance.

Proposition 1.2. *A graph G is a Boolean sum of κ cliques if and only if G has a Boolean representation in a set of cardinality κ.*

Proof: If G is the Boolean sum of a family $(C_i)_{i\in I}$ of κ cliques, then let $f : V(G) \to \wp(I)$ defined by setting $f(x) := \{i \in I : x \in C_i\}$. This defines a Boolean representation in I. Conversely, if $f : V(G) \to$

$\wp(E)$ is a Boolean representation in a set E, then set $C_i := \{x \in V(G) : i \in f(x)\}$ for $i \in E$. Then G is the Boolean sum of the family $(G_i)_{i \in E}$, where $G_i := (V(G), [C_i]^2)$ for each $i \in E$. □

We note that the Boolean dimension of a graph and of the graph obtained by removing some isolated vertices are the same. Hence $\dim_{\text{Bool}}(G) = 1$ if and only if it is of the form $G = K_X^V$ with $|X| \geq 2$. Since every graph $G := (V, E)$ can be viewed as the Boolean sum of its edges, the Boolean dimension of G is always defined, and is at most the number of edges, that is, at most the cardinality $|[V]^2|$ of $[V]^2$. If V is infinite, then $|[V]^2| = |V|$; hence $\dim_{\text{Bool}}(G) \leq |V|$ (but see Question 1.1 below). If V is finite, with n elements, then $\dim_{\text{Bool}}(G) \leq n - 1$ (see Belkhechine et al. (2010). By induction on n: pick $x \in V$ and observe that $G = (V, E \setminus \{e \in E : x \in e\}) \dot{+} K^V_{N_G(x)} \dot{+} K^V_{N_G(x) \cup \{x\}})$. In fact, paths on n vertices are the only n-vertex graphs with Boolean dimension $n - 1$, see Theorem 2.4, a result that requires some ingredients developed below.

Recall that a *module* in a graph G is any subset A of $V(G)$ such that for every $a, a' \in A$ and $b \in V(G) \setminus A$, we have $a \sim b$ if and only if $a' \sim b$. A *duo* is any two-element module (e.g., see Courcelle and Delhommé (2008) for an account of the modular decomposition of graphs).

Lemma 1.3. *If a graph G has no duo then every Boolean representation is one to one. In particular,* $\dim_{\text{Bool}}(G) \geq \log_2(|V(G)|)$.

Proof: Observe that if f is a representation and v is in the range of f, then $f^{-1}(v)$ is a module and this module is either a clique or an independent set. □

Question 1.1. The inequality in Lemma 1.3 may be strict if $V(G)$ is finite. Does $\dim_{\text{Bool}}(G) \leq \log_2(|V(G)|)$ when $V(G)$ is infinite? The answer may depend on some set theoretical hypothesis (see Example 1.9). But we do not known if the Boolean dimension of every graph on at most a continuum of vertices is at most countable. Same question may be considered for trees.

Let E be a set; denote by $O(E)^{\neg\perp}$ the graph whose vertices are the subsets of E, two vertices X and Y being linked by an edge if they are distinct and their intersection is finite and odd. If κ is a cardinal, we set $O(\kappa)^{\neg\perp}$ for any graph isomorphic to $O(E)^{\neg\perp}$, where E is a set of cardinality κ.

Theorem 1.4. *A graph G with no duo has Boolean dimension at most κ if and only if it is embeddable in $O(\kappa)^{\neg\perp}$. The Boolean dimension of $O(\kappa)^{\neg\perp}$ is at most κ. It is equal to κ if $\kappa \geq 2$ and κ is at most countable, or if κ is uncountable and (GCH) holds.*

Proof: If there is an embedding f from G in a graph of the form $O(E)^{\neg\perp}$, then f is a Boolean representation of G, hence $\dim_{\text{Bool}}(G) \leq |E|$. Conversely, if G has no duo and has a Boolean representation f in a set E then, by Lemma 1.3, f is an embedding of G in $O(E)^{\neg\perp}$. Let E be a set of cardinality κ. For each $X \in V(O(E)^{\neg\perp}) = \wp(E)$ set $f(X) := X$ viewed as a subset of $\wp(E)$. The map f is a Boolean representation, hence $\dim_{\text{Bool}} O(E)^{\neg\perp} \leq \kappa$. Alternatively, set $C_i := \{X \in \wp(E) : i \in X\}$ for each $i \in E$. Then $O(E)^{\neg\perp}$ is the Boolean sum of the $[C_i]^2, i \in E$. If $\kappa = 2$, a simple inspection shows that the Boolean dimension of $O(E)^{\neg\perp}$ is κ. If $\kappa \geq 3$ then $O(E)^{\neg\perp}$ has no duo. This relies on the following claim.

Claim 1.5. *If A, B are two distinct subsets of E, then there is a subset C of E, distinct from A and B, with at most two elements such that the cardinalities of the sets $A \cap C$ and $B \cap C$ cannot have the same parity.*

Indeed, we may suppose that $A \nsubseteq B$. Pick $x \in A \setminus B$. If $|A| > 1$, then set $C := \{x\}$. If not, then $A = \{x\}$. In this case, either B is empty and $C := \{x, y\}$, with $y \neq x$ will do, or B is nonempty, in which case, we may set $C := \{y\}$, where $y \in B$ if $|B| > 1$, or $C := \{y, z\}$, where $B = \{y\}$ and $z \in E \setminus (A \cup B)$.

Since $O(E)^{\neg\perp}$ has no duo, Lemma 1.3 ensures that $\dim_{\mathrm{Bool}}(O(E)^{\neg\perp}) \geq \log_2(|V(O(E))|) = \log_2(2^\kappa)$. If κ is at most countable, or κ is uncountable and (GCH) holds, then this last quantity is κ. This completes the proof of the theorem. □

We can obtain the same conclusion with a weaker hypothesis than (GCH).

Lemma 1.6. *Let κ be an infinite cardinal. If $\mu^\omega < \kappa$ for every $\mu < \kappa$, then* $\dim_{\mathrm{Bool}}(O(\kappa)^{\neg\perp}) = \kappa$.

Proof: The proof relies on the following claim, which is of independent interest.

Claim 1.7. *Let μ^ω be the cardinality of the set of countable subsets of an infinite cardinal μ. Then the cliques in $O(\mu)^{\neg\perp}$ have cardinality at most μ^ω.*

The proof relies on a property of almost disjoint families. Let us recall that an *almost disjoint family* is a family $\mathcal{A} := (A_\alpha)_{\alpha \in I}$ of sets such that the intersection $A_\alpha \cap A_\beta$ is finite for $\alpha \neq \beta$. Note that if $\mathcal{C}$ is a clique in $O(\mu)^{\neg\perp}$, then for every pair of distinct sets X, Y in $\mathcal{C}$, the intersection $X \cap Y$ is finite and its cardinality is odd. Hence, $\mathcal{C}$ is an almost disjoint family.

To prove our claim it suffices to prove the following claim, well known by set theorists.

Claim 1.8. *There is no almost disjoint family of more that μ^ω subsets of an infinite set of cardinality μ.*

Proof of Claim 1.8. Suppose that such a family $\mathcal{A} := (A_\alpha)_{\alpha \in I}$ exists, with $|I| > \mu^\omega$. Since $\mu^{<\omega} = \mu$, we may suppose that each A_α is infinite and then select a countable subset B_α of A_α. The family $\mathcal{B} := (B_\alpha)_{\alpha \in I}$ is almost disjoint, but since $|I| > \mu^\omega$, there are $\alpha \neq \beta$ such that $B_\alpha = B_\beta$, hence $B_\alpha \cap B_\beta$ is infinite, contradicting the fact that $\mathcal{B}$ is an almost disjoint family. □

Now the proof of the lemma goes as follows. Suppose that $\dim_{\mathrm{Bool}}(O(\kappa)^{\neg\perp} = \mu < \kappa$. Then there is an embedding from the graph $O(\kappa)^{\neg\perp}$ into the graph $O(\mu)^{\neg\perp}$. Trivially, $O(\kappa)^{\neg\perp}$ contains cliques of cardinality at least κ. Hence $O(\mu)^{\neg\perp}$ too. But since $\mu^\omega < \kappa$, Claim 1.7 says that this is impossible. Thus $\dim_{\mathrm{Bool}}(O(\kappa)^{\neg\perp}) = \kappa$. □

We thank Avraham (2021) for providing Claim 1.8.

Examples 1.9. For a simple illustration of Lemma 1.6, take $\kappa = (2^{\aleph_0})^+$ the successor of $2^{\aleph_0}$. For an example, negating (GCH), suppose $\omega_1 = 2^{\aleph_0}$, $\kappa = \omega_2$, $\omega_3 = 2^{\omega_1} = 2^{\omega_2}$. In this case, $\dim_{\mathrm{Bool}}(O(\kappa)^{\neg\perp}) = \kappa$ and $\log_2(2^\kappa) = \omega_1 < \kappa$.

Question 1.2. Does the equality $\dim_{\mathrm{Bool}}(O(\kappa)^{\neg\perp}) = \kappa$ hold without any set theoretical hypothesis?

Remark 1.10. Theorem 1.4 asserts that $O(\kappa)^{\neg\perp}$ is universal among graphs with no duo of Boolean dimension at most κ (that is embeds all graphs with no duo of dimension at most κ), but we do not know which graphs on at most 2^κ vertices embed in $O(\kappa)^{\neg\perp}$.

In contrast with Claim 1.7 we have:

Lemma 1.11. *For an infinite cardinal κ, the graph $O(\kappa)^{\neg\perp}$ embeds a graph made of 2^κ disjoint edges. It embeds also some trees made of 2^κ vertices.*

Proof: Let G be the graph made of 2^κ disjoint edges $\{a_\alpha, b_\alpha\}$ with $\alpha \in 2^\kappa$. We show that G is isomorphic to an induced subgraph of $O(E)^{\neg\perp}$, where E is the set $[\kappa]^{<\omega}$ of finite subsets of κ, augmented of an extra element r. Since $|E| = \kappa$, this proves our first statement. For the purpose of the proof, select 2^κ subsets X_α of κ which are pairwise incomparable with respect to inclusion and contain an infinite subset X. For each $\alpha \in 2^\kappa$, let $A_\alpha := [X_\alpha]^{<\omega} \cup \{r\}$ and $B_\alpha := E \setminus [X_\alpha]^{<\omega}$. We claim that the subgraph H of $O(E)^{\neg\perp}$ induced by $\{A_\alpha : \alpha \in 2^\kappa\} \cup \{B_\alpha : \alpha \in 2^\kappa\}$ is a direct sum of the edges $\{A_\alpha, B_\alpha\}$ for $\alpha \in 2^\kappa$. That A_α and B_α form an edge is obvious: their intersection is the one element set $\{r\}$. Now, let $\alpha \neq \beta$. We claim that the three intersections $A_\alpha \cap A_\beta$, $A_\alpha \cap B_\beta$ and $B_\alpha \cap B_\beta$ are all infinite. For the first one, this is obvious (it contains $[X_\alpha \cap X_\beta]^{<\omega}$), for the next two, use the fact that the A_α are up-directed with respect to inclusion, hence the difference $A_\alpha \setminus A_\beta$ is cofinal in A_α, thus must be infinite, and the union $A_\alpha \cup A_\beta$ cannot cover $[\kappa]^{<\omega}$, hence its complement is infinite. It follows that the graph H contains no other edges than the pairs $\{A_\alpha, B_\alpha\}$'s. This proves that H is isomorphic to G, and yields our first statement. For the second statement, add $R := [X]^{<\omega} \cup \{r\}$ to the set of vertices of H. We get a tree. Indeed, for each α, the vertices R and A_α do not form an edge in $O(E)^{\neg\perp}$ (indeed, $R \cap A_\alpha = [X]^{<\omega}$ hence is infinite), while for each β, the vertices R and A_β form an edge (since $R \cap A_\beta = \{r\}$). □

For infinite graphs with finite Boolean dimension, a straightforward application of Tychonoff's theorem yields the following result.

Theorem 1.12. *Let $n \in \mathbb{N}$. For every graph G, $\dim_{\mathrm{Bool}}(G) \leq n$ if and only if $\dim_{\mathrm{Bool}}(G_{\restriction X}) \leq n$ for every finite subset X of $V(G)$.*

Proof: Suppose that the second condition holds. For every finite subset X of $V(G)$ let U_X be the set of maps f from $V(G)$ into the powerset $K := \wp(\{1, \dots, n\})$ such that the restriction $f_{\restriction X}$ is a Boolean representation of $G_{\restriction X}$ in $\{1, \dots, n\}$. Each such set U_X is nonempty and closed in the set $K^{V(G)}$ equipped with the product topology, the set K being equipped with the discrete topology. Every finite intersection $U_{X_1} \cap \dots \cap U_{X_\ell}$ contains $U_{X_1 \cup \dots \cup X_\ell}$ hence is nonempty. The compactness of $K^{V(G)}$ ensures that the intersection of all of those sets is nonempty. Any map in this intersection is a Boolean representation of G. □

Examples of graphs with finite Boolean dimension are given at the end of the next subsection.

1.2 Geometric notions of dimensions of graphs

We introduce three notions of dimensions: geometric, inner, and symplectic, all based on bilinear forms. We prove that if the Boolean dimension of a graph is finite, then it coincides with the inner dimension, and either these dimensions minus 1 coincide with the geometric and the symplectic dimension, or they coincide with the geometric dimension, the symplectic being possibly larger (Theorem 1.18). We note before all that in general, the Boolean dimension is not based on a bilinear form. It uses the map $\varphi : \wp(E) \to 2 := \{0, 1\}$ defined by setting $\varphi(X, Y) := 1$ if $|X \cap Y|$ is finite and odd and 0 otherwise. But except when E is finite, it is not bilinear on $\wp(E)$ equipped with the symmetric difference.

Let $\mathbb{F}$ be a field, and let U be a vector space over $\mathbb{F}$, and let φ be a bilinear form over U. We recall that this form is *symmetric* if $\varphi(x, y) = \varphi(y, x)$ for all $x, y \in U$. Two vectors x, y are *orthogonal* if $\varphi(x, y) = 0$. A vector $x \in U$ is *isotropic* if $\varphi(x, x) = 0$. The *orthogonal* of a subset X of U is the subspace $X^\perp := \{y \in U : \varphi(x, y) = 0 \text{ for all } x \in X\}$. We set $x^\perp$ instead of $\{x\}^\perp$. We recall that φ is *degenerate* if there is some $x \in U \setminus \{0\}$ such that $\varphi(x, y) = 0$ for all $y \in U$. The form φ is said to be *alternating* if each $x \in U$ is isotropic, in which case (U, φ) is called a *symplectic space*. The form φ is an *inner form* or a *scalar product* if U has an *orthonormal basis* (made of non-isotropic and pairwise othogonal vectors).

Definition 1.13. Let U be a vector space equipped with a symmetric bilinear form φ. Let G be a graph. We say that a map $f\colon V(G) \to U$ is a *geometric representation* of G in (U,φ) if for all $u, v \in V(G), u \neq v$, we have $u \sim v$ if and only if $\varphi(f(u), f(v)) \neq 0$. The *geometric dimension* of G, denoted by $\dim_{\mathrm{geom}}(G)$, is the least cardinal κ for which there exists a geometric representation of G in a vector space U of dimension κ equipped with a symmetric bilinear form φ. The *symplectic dimension* of G, denoted by $\dim_{\mathrm{symp}}(G)$, is the least cardinal κ for which there exists a symplectic space (U,φ) in which G has a geometric representation. The *inner dimension* of G, denoted by $\dim_{\mathrm{inn}}(G)$, is the least cardinal κ for which G has a geometric representation in a vector space of dimension κ equipped with a scalar product.

The notions of geometric and symplectic dimension were considered by several authors, for example, Garzon (1987); Godsil and Royle (2001a). There is an extensive literature about this subject (e.g. Fallat and Hogben (2007); Grout (2010)), and notably the role of the field. But apparently, the Boolean dimension was not considered.

Except in subsection 1.4, we consider these notions only for the 2-element field $\mathbb{F}_2$, identified with the set $\{0,1\}$. If U has finite dimension, say k, we identify it with $\mathbb{F}_2^k$, the set of all k-tuples over $\{0,1\}$; the basis $(e_i)_{i:=1,\dots,k}$, where e_i is the 0-1-vector with a 1 in the i-th position and 0 elsewhere, is orthonormal; the scalar product of two vectors $x := (x_1,\dots,x_k)$ and $y := (y_1,\dots,y_k)$ of $\mathbb{F}_2^k$ is then $\langle x \mid y\rangle := x_1y_1 + \cdots + x_ky_k$. We recall the following dichotomy result.

Theorem 1.14. *A nondegenerate bilinear symmetric form φ on a finite k-dimensional space U over the two-element field $\mathbb{F}_2$ falls into two types. Either φ is non-alternating and (U,φ) is isomorphic to $(\mathbb{F}_2^k, |)$ with the scalar product, or φ is alternating, k is even, and (U,φ) is isomorphic to the symplectic space $H(k) := (\mathbf{1}^{\neg\perp}, |_{\restriction \mathbf{1}^{\neg\perp}})$, where $\mathbf{1}^{\neg\perp}$ is the orthogonal of $\mathbf{1} := (1,\dots,1)$ with respect to the scalar product $|$ on $\mathbb{F}_2^{k+1}$.*

For reader's convenience, we give a proof. The proof, suggested by Christian Delhommé, is based on two results exposed in Algebra, Vol. 3, of Cohn (1991). Let (U,φ) be as stated in the above theorem. Case 1: φ is not symplectic, that is $\varphi(x,x) \neq 0$ for some vector x. We apply Proposition 7.1 page 344 of Cohn (1991), namely: *If U is a vector space of characteristic 2 and φ is a symmetric bilinear form which is not alternating, then U has a orthogonal basis*. Since φ is nondegenerate and the field if $\mathbb{F}_2$, any orthogonal basis is orthonormal, hence φ is a scalar product. Case 2: φ is symplectic. In this case, Lemma 5.1, p.331 of Cohn (1991) asserts in particular that: *Every symplectic space, (that is a space equipped with a bilinear symmetric form which is nondegenerate and alternating) on an arbitrary field is a sum of hyperbolic planes*. Thus k is even and in our case U is isomorphic to any symplectic space with the same dimension, in particular to $H(k)$.

When dealing with these notions of dimension, we may always consider nondegenerate forms, hence in the case of finite dimensional representation, Theorem 1.14 applies. In fact Lemma 1.3 and Theorem 1.4 extend.

Let U be a vector space over $\mathbb{F}_2$ and φ a symmetric bilinear form defined on U with values in $\mathbb{F}_2$. Let $O_\varphi^{\neg\perp}$ be the graph of the non-orthogonality relation on U, that is, the graph whose edges are the pairs of distinct elements x and y such that $\varphi(x,y) = 1$. If k is an integer, then we denote by $O_{\mathbb{F}_2}^{\neg\perp}(k)$ the graph on $\mathbb{F}_2^k$ of the non-orthogonality relation associated with the inner product $|$. Similarly, for k even, let $O_H^{\neg\perp}(k)$ be the graph on $H(k)$, the orthogonal of $\mathbf{1} := (1,\dots,1)$ with respect to the scalar product $|$ on $(\mathbb{F}_2)^{k+1}$, equipped with the symplectic form induced by the scalar product.

Lemma 1.15. *If $\dim(U)$, the dimension of the vector space U, is at least 3, then the graph $O_\varphi^{\neg\perp}$ has no duo if and only if φ is nondegenerate. Hence, $\dim_{\mathrm{geom}}(O_\varphi^{\neg\perp}) = \dim(U)$ when φ is nondegenerate.*

Proof: Suppose that φ is degenerate. Pick a nonzero element a in the kernel of φ. Then, as it is easy to check, the 2-element set $\{0, a\}$ is a module of $O_{\varphi}^{\neg\perp}$. Conversely, let $\{a, b\}$ be a duo of $O_{\varphi}^{\neg\perp}$. We claim that $c := a+b$ belongs to the kernel of φ, that is $\varphi(x, c) = 0$ for every $x \in U$. Indeed, if $x \notin \{a, b\}$, then $\varphi(x, a) = \varphi(x, b)$, hence $\varphi(x, c) = 0$ since $\{a, b\}$ is a module. If $x \in \{a, b\}$ (e.g. $x := a$), then since $\dim(U) \geq 3$, we may pick some $z \notin \operatorname{span}\{a, b\} := \{0, a, b, a+b\}$, hence $\varphi(z, c) = 0$. Since $z + a \notin \{a, b\}$, $\varphi(z + a, c) = 0$. It follows that $\varphi(a, c) = 0$, proving our claim. According to Lemma 1.3, every representation of $O_{\varphi}^{\neg\perp}$ is one to one; since the identity map is a representation, we have $\dim_{\mathrm{geom}}(O_{\varphi}^{\neg\perp}) = \dim(U)$. □

We give below an existential result. The proof of the second item is based on the Δ-system lemma (see (Kunen, 2011; Rinot) for an elementary proof) that we recall now.

Lemma 1.16. *Suppose that κ is a regular uncountable cardinal, and $\mathcal{A} := (A_\alpha)_{\alpha\in\kappa}$ is a family of finite sets. Then there exist a subfamily $\mathcal{B} := (A_\alpha)_{\alpha\in K}$, where the cardinality of K is κ, and a finite set R such that $A_\alpha \cap A_\beta = R$ for all distinct $\alpha, \beta \in K$.*

Theorem 1.17.

1. *Every graph has a symplectic dimension, and hence, it has a geometric one. However:*
2. *not every graph has an inner dimension, e.g., a graph with κ vertices, with κ regular, and no clique and no independent set of κ vertices, does not have an inner representation; on an other hand:*
3. *every locally finite graph has an inner dimension.*

Proof:

1. Let G be a graph, and $\kappa := |V(G)|$. Let U be a vector space over $\mathbb{F}_2$ with dimension κ (e.g., $U := \mathbb{F}_2^{[V(G)]}$, the set of maps $f : V(G) \to \mathbb{F}_2$ which are 0 almost everywhere). Define a symplectic form φ on a basis $\mathcal{B} := \{b_v : v \in V(G)\}$ of U indexed by the elements of $V(G)$ (e.g., b_v is the map from $V(G)$ to $\mathbb{F}_2$ defined by $b_v(v) = 1$ and $b_v(u) = 0$ for $u \neq v$). For that, set $\varphi(b_u, b_v) := 1$ if $u \neq v$ and $u \sim v$; in particular $\varphi(b_v, b_v) = 0$ for every $v \in V(G)$. Then extend φ on U by bilinearity. Since the vectors of the basis are isotropic and $\mathbb{F}_2$ has characteristic two, φ is symplectic. By construction, the map $v \to b_v$ is a representation of G in (U, φ). Hence G has a symplectic dimension.
2. An inner representation of a graph G reduces to a map f from $V(G)$ into the vector space $[E]^{<\omega}$ of finite subsets of a set E equipped with the symmetric difference such that for every two-element subset $e := \{u, v\}$ of $V(G)$, we have $e \in E(G)$ if and only if $|f(u) \cap f(v)|$ is odd. Suppose that $V(G) = \kappa$ and no subset of $V(G)$ of cardinality κ is a clique or an independent set. According to Ramsey's theorem, κ is uncountable. Apply Lemma 1.16 to $\mathcal{A} := (f(u))_{u\in V(G)}$. Let $\mathcal{B} := (f(u))_{u\in K}$ be a subfamily of $\mathcal{A}$, where K has cardinal κ, and let R be given by this lemma. Since $f(u) \cap f(v) = R$ for all every $u, v \in K$, the set K is a clique or an independent set depending on the fact that the cardinality of R is odd or even. Hence, if G has no clique and no independent set of κ vertices, it cannot have an inner representation. A basic example on cardinality $\aleph_1$ is provided by the comparability graph G of a Sierpinskization of a subchain A of the reals of cardinality $\aleph_1$ with an order of type ω_1 on A.
3. Let $E := E(G)$. Let $[E]^{<\omega}$ be the collection of finite subsets of E; equipped with the symmetric difference Δ, $[E]^{<\omega}$ is a vector space over $\mathbb{F}_2$; the one-element subsets of E form a basis; the map $\varphi : [E]^{<\omega} \times [E]^{<\omega} \to \mathbb{F}_2$ defined by setting $\varphi(X, Y) = 1$ if $|X \cap Y|$ is odd and $\varphi(X, Y) = 0$

otherwise is a bilinear form for which the one-element subsets of E form an orthonormal basis. Hence φ is an inner product. Let $f : V(G) \to \wp(E)$ be defined by setting $f(x) := E_G(x)$ $(= \{e \in E : x \in e\})$. Since for any pair of distinct vertices $x, y \in V(G)$, $|E_G(x) \cap E_G(y)| = 1$ amounts to $\varphi(f(x), f(y)) = 1$, f is an inner representation of G.

□

As noted by Delhommé (2021), the Boolean dimension can be strictly smaller than the geometric dimension. For an example, if κ is an infinite cardinal, the geometric dimension of $O(\kappa)^{\neg\perp}$ is 2^κ while its Boolean dimension is at most κ. Indeed, from Theorem 1.17, $O(\kappa)^{\neg\perp}$ has a geometric representation in a vector space U. As for any representation, Lemma 1.3 is still valid; since $O(\kappa)^{\neg\perp}$ has no duo (for $\kappa \geq 3$) the cardinality of U is at least 2^κ, thus the dimension of the vector space U is 2^κ, while $O(\kappa)^{\neg\perp}$ has a Boolean representation in a set of cardinality κ.

Problem 1.3. Does every countable graph has an inner dimension? [i]

1.3 Graphs with finite geometric dimension

Theorem 1.18. *If the Boolean dimension of a graph G is finite, then it is equal to the inner dimension of G and either*

1. *the geometric dimension, the symplectic dimension and the Boolean dimension of G are equal,*

   or

2. *the geometric dimension and the symplectic dimension of G are equal to the Boolean dimension of G minus* 1*,*

   or

3. *the geometric dimension and the Boolean dimension of G are equal and are strictly less than the symplectic dimension of G, in which case the difference between these two numbers can be arbitrarily large.*

Proof: The first assertion is obvious. By definition, $\dim_{\text{geom}}(G) \leq \min\{\dim_{\text{Bool}}(G), \dim_{\text{symp}}(G)\}$. Apply Theorem 1.14. Let $k := \dim_{\text{geom}}(G)$. If $k \neq \dim_{\text{Bool}}(G)$, then G is representable into $H(k)$ and thus in $\mathbb{F}_2^{k+1}$, hence (2) holds. If $k = \dim_{\text{Bool}}(G)$, then $\dim_{\text{symp}}(G) \geq k$. The examples given in (a) below show that the difference $\dim_{\text{symp}}(G) - \dim_{\text{Bool}}(G)$ can be large. □

We give some examples when the graphs are finite.

Examples 1.19. (a) $\dim_{\text{geom}}(K_X^V) = \dim_{\text{Bool}}(K_X^V) = 1$ if $|X| \geq 2$ (in fact they equal 0 if $|X| \leq 1$) and $\dim_{\text{symp}}(K_X^V) = 2k$ if $|X| \in \{2k, 2k+1\}$.

(b) $\dim_{\text{geom}}(O_{\mathbb{F}_2}^{\neg\perp}(k)) = \dim_{\text{Bool}}(O_{\mathbb{F}_2}^{\neg\perp}(k)) = k$ for $k \geq 2$, and 0 otherwise.

(c) $\dim_{\text{geom}}(O_H^{\neg\perp}(k)) = \dim_{\text{symp}}(O_H^{\neg\perp}(k)) = \dim_{\text{Bool}}(O_H^{\neg\perp}(k)) - 1 = k$ for $k = 2m \geq 4$, and $\dim_{\text{geom}}(O_H^{\neg\perp}(2)) = \dim_{\text{Bool}}(O_H^{\neg\perp}(2)) = \dim_{\text{symp}}(O_H^{\neg\perp}(2)) - 1 = 1$.

---

[i] Norbert Sauer informed us on january 2022 that the answer is positive

These examples are extracted from Belkhechine et al. (2012). The paper being unpublished, we give a hint below. We use the following lemma.

Lemma 1.20. *If $G := (V, E)$ is a graph for which $\dim_{\mathrm{symp}}(G) = 2k \in \mathbb{N}$, then every clique of G has at most $2k + 1$ elements.*

This fact is a straightforward consequence of the following claim which appears equivalently formulated in van Lint and Wilson (2001) as Problem 19O.(i), page 238.

Claim 1.21. *If $\ell + 1$ subsets A_i, $i < \ell + 1$, of an ℓ-element set A have odd size, then there are $i, j < \ell + 1$, $i \neq j$ such that $A_i \cap A_j$ has odd size.*

We prove now that the examples satisfy the stated conditions.

Item (a). The first part is obvious. For the second part, we use Claim 1.21 and Lemma 1.20. Indeed, let $f : V(G) \to H(2k)$. Composing with the involution h of $\mathbb{F}_2^{2k+1}$ we get a representation in $\mathbf{1} + H(2k)$, where the involution h is defined by $h(x) = x + \mathbf{1}$, where $\mathbf{1} := (1, 1, \dots, 1) \in F_2^{2k+1}$. The image of a clique of G yields subsets of odd size such that the intersection of distinct subsets has even size. Thus from Claim 1.21 above there are no more than $2k + 1$ such sets.

With that in hand, we prove the desired equality $\dim_{\mathrm{symp}}(K_X^{V(G)}) = 2k$ if $|X| \in \{2k, 2k + 1\}$.

Indeed, let X be an n-element subset of $V(G)$ and let $(x_i)_{i<n}$ be an enumeration of X. Let k with $n \leq 2k + 1$ and $f : V(G) \to \mathbb{F}_2^{2k+1}$ be defined by $f(x) = 0$ if $x \in V(G) \setminus X$ and $f(x) := (b_j)_{j<2k+1}$, where $b_j = 1$ for all $j \neq i$ and $b_i = 0$ if $x = x_i$. Clearly, f is a representation of G in $H(2k)$, thus $\dim_{\mathrm{symp}}(K_X^{V(G)}) \leq 2k$. The reverse inequality follows from Lemma 1.20.

Item (b). If $k = 1$, the graph $O_{\mathbb{F}_2}^{-\perp}(k)$ is made of two isolated vertices, and if $k = 2$ the graph is a path on three vertices plus an isolated vertex, their respective Boolean dimensions are 1 and 2, as claimed. If $k \geq 3$ the result follows from the conclusion of Lemma 1.15.

Item (c) If $k = 2$, the graph $O_H^{-\perp}(k)$ is made of a clique on three vertices plus an isolated vertex, hence its Boolean dimension is 1. If $k \geq 4$, the equality $\dim_{\mathrm{geom}}(O_H^{-\perp}(k)) = \dim_{\mathrm{symp}}(O_H^{-\perp}(k))$ follows from the conclusion of Lemma 1.15. The number of edges of $O_H^{-\perp}(k)$ and $O_{\mathbb{F}_2}^{-\perp}(k)$ are different, hence $O_H^{-\perp}(k)$ cannot have a Boolean representation in $(\mathbb{F}_2^k, |)$. Since it has a representation in $(\mathbb{F}_2^{k+1}, |)$, the result follows. □

The paper by Godsil and Royle (2001a) contains many more results on the symplectic dimension over $\mathbb{F}_2$ of finite graphs.

1.4 Dimension and rank

We compute the symplectic dimension and the geometric dimension of a graph G in terms of its adjacency matrix.

Let $n \in \mathbb{N}$. Let A be an $n \times n$ symmetric matrix with coefficients in a field $\mathbb{F}$. We denote by $\mathrm{rank}_{\mathbb{F}}(A)$ the rank of A computed over the field $\mathbb{F}$. The *minrank* of A, denoted by $\mathrm{minrank}_{\mathbb{F}}(A)$, is the minimum of $\mathrm{rank}_{\mathbb{F}}(A + D)$, where D is any diagonal symmetric matrix with coefficients in $\mathbb{F}$. If $\mathbb{F} = \mathbb{F}_2$, we denote these quantities by $\mathrm{rank}_2(A)$ and $\mathrm{minrank}_2(A)$. Let $G := (V, E)$ be a graph on n vertices. Let $v_1, \dots, v_n$ be an enumeration of V. The *adjacency matrix* of G is the $n \times n$ matrix $A(G) := (a_{i,j})_{1 \leq i,j \leq n}$ such that $a_{i,j} = 1$ if $v_i \sim v_j$ and $a_{i,j} = 0$ otherwise.

Theorem 1.22. *If G is a graph on n vertices, then the symplectic and the geometric dimensions of G over a field $\mathbb{F}$ are respectively equal to the rank and the minrank of $A(G)$ over $\mathbb{F}$.*

An $n \times n$ symmetric matrix $B := (b_{i,j})_{1 \le i,j \le n}$ over a field $\mathbb{F}$ is *representable* as the matrix of a symmetric bilinear form φ on a vector space U over a field $\mathbb{F}$ if there exists n vectors $u_1, \dots, u_n$ in U, not necessarily distinct, such that $b_{i,j} = \varphi(u_i, u_j)$ for all $1 \le i, j \le n$.

The matrix B can be represented in $U := \mathbb{F}^n$, where $(u_i)_{1 \le i \le n}$ is the canonical basis and $\varphi(u_i, u_j) = b_{i,j}$. According to the following lemma (see Corollary 8.9.2 p. 179 of Godsil and Royle (2001b)), there is a representation in a vector space whose dimension is the rank of the matrix B.

Lemma 1.23. *An $n \times n$ symmetric matrix B of rank r has a principal $r \times r$ submatrix of full rank.*

The following result shows that this value is optimum.

Lemma 1.24. *The smallest dimension of a vector space in which a symmetric matrix B is representable is the rank of B.*

Proof: It is an immediate consequence of the following facts, whose proofs are a simple exercise in linear algebra.

1) Let $r := \operatorname{rank}(B)$. Then $r \le \dim(U)$ for any vector space U in which B is representable. Let φ be a bilinear form on U, and let $u_1, \dots, u_n$ be n vectors of U such that $\varphi(u_i, u_j) = b_{ij}$ for all $1 \le i, j \le n$, where $(b_{i,j})_{1 \le i,j \le n} = B$. Let $B(j_1), \dots, B(j_r)$ be r linearly independent column vectors of B with indices $j_1, \dots, j_r$. We claim that the corresponding vectors $u_{j_1}, \dots, u_{j_r}$ are linearly independent in U. Suppose that a linear combination $\sum_{k=1}^{r} \lambda_{j_k} u_{j_k}$ is zero. Then, for every vector $u \in U$, $\varphi(\sum_{k=1}^{r} \lambda_{j_k} u_{j_k}, u) = 0$. This rewrites as $\sum_{k=1}^{r} \lambda_{j_k} \varphi(u_{j_k}, u) = 0$. In particular, $\sum_{k=1}^{r} \lambda_{j_k} \varphi(u_{j_k}, u_i) = 0$ for every $i = 1, \dots, n$. That is, $\sum_{k=1}^{r} \lambda_{j_k} B_{j_k} = 0$. Since these column vectors are linearly independent, the λ_{j_k}'s are zero. This proves our claim.

2) Suppose that φ is nondegenerate and U is spanned by the vectors $u_1, \dots, u_n$. Then $r \ge \dim(U)$. The proof follows the same lines as above. Let $s := \dim(U)$. Then, among the u_j's there are s linearly independent vectors, say $u_{j_1}, \dots, u_{j_s}$. We claim that the column vectors $B(j_1), \dots, B(j_s)$ are linearly independent. Suppose that a linear combination $\sum_{k=1}^{s} \lambda_k B_{j_k}$ is zero. This yields $\sum_{k=1}^{s} \lambda_k \varphi(u_{j_k}, u_i) = 0$ for every i, $1 \le i \le n$, hence $\varphi(\sum_{k=1}^{s} \lambda_k u_{j_k}, u_i) = 0$. Since the u_i's generate U, we have $\varphi(\sum_{k=1}^{s} \lambda_k u_{j_k}, u) = 0$ for every $u \in U$. Since the form φ is nondegenerate, $\sum_{k=1}^{s} \lambda_k u_{j_k} = 0$. Since the vectors $u_{j_1}, \dots, u_{j_s}$ are linearly independent, the λ_k's are all zero. This proves our claim.

3) Suppose that B is representable in a vector space U equipped with a symmetric bilinear form φ. Then B is representable in a quotient of U equipped with a nondegenerate bilinear form. □

Theorem 1.22 follows immediately from Lemma 1.24.

Remark 1.25. Theorem 1.22 for the symplectic dimension of graphs over $\mathbb{F}_2$ is due to Godsil and Royle (2001a). The minrank over several fields has been intensively studied, see Fallat and Hogben (2007) for a survey. These authors consider the problem of minrank of graphs, and obtain a combinatorial description for the minimum rank of trees. In the next section, we only state that in case of trees, the Boolean dimension, geometric dimension and the minimum rank coincide, thus the formula given in Theorem 2.9 below for the Boolean dimension gives yet another combinatorial description for the minimum rank of a tree.

2 Boolean dimension of trees

In this section, we show that there is a nice combinatorial interpretation for the Boolean dimension of trees. We mention first the following result of Houmem Belkhechine et al. Belkhechine et al. (2012).

Lemma 2.1. *Let $G := (V, E)$ be a graph, with $V \neq \emptyset$. Let $m \in \mathbb{N}$, and let $f\colon V \to \mathbb{F}_2^m$ be a representation of G in the vector space $\mathbb{F}_2^m$ equipped with a symmetric bilinear form φ. Let $A \subseteq V$ such that $A \neq \emptyset$. Suppose that for all finite $X \subseteq A, X \neq \emptyset$, there exists $v \in V \setminus X$ such that $|N_G(v) \cap X|$ is odd. Then $\{f(x) \mid x \in A\}$ is linearly independent in the vector space $\mathbb{F}_2^m$.*

Proof: Let X be a non empty finite subset of A. We claim that $\sum_{x\in X} f(x) \neq 0$. Indeed, let $v \in V \setminus X$ such that $|N_G(v) \cap X|$ is odd. We have $\varphi(\sum_{x\in X} f(x), f(v)) = \sum_{x\in X} \varphi(f(x), f(v))$. This sum is equal to $|N_G(v) \cap X|$ modulo 2. Thus $\sum_{x\in X} f(x) \neq 0$ as claimed. Since this holds for every finite subset X of A, the conclusion follows. □

This suggests the following definition.

Definition 2.2 ((Belkhechine et al., 2012)). Let $G := (V, E)$ be a graph. A set $A \subset V$ is called *independent* $\pmod 2$ if for all finite $X \subseteq A, X \neq \emptyset$, there exists $v \in V \setminus X$ such that $|N_G(v) \cap X|$ is odd, otherwise A is said to be dependent $\pmod 2$. Let $\mathrm{ind}_2(G)$ be the maximum size of an independent set $\pmod 2$ in G. **From now, we omit $\pmod 2$ unless it is necessary to talk about independence in the graph theoretic sense.**

Corollary 2.3. *For every graph G, we have* $\mathrm{ind}_2(G) \leq \dim_{\mathrm{geom}}(G)$.

Problem 2.1. Does the equality hold?

Note that the independent sets $\pmod 2$ of a graph do not form a matroid in general. Indeed, let G be made of six vertices, three, say $\{a, b, c\}$ forming a clique, the three others, say a', b', c' being respectively connected to a, b and c. Then $\{a', a, b, c\}$ is independent $\pmod 2$, hence $4 \leq \mathrm{ind}_2(G)$. Also, $\{a', b', c'\}$ is independent $\pmod 2$ but cannot be extended to a larger independent set $\pmod 2$. Since G is the Boolean sum of a 3-vertex clique and three edges, $\dim_{\mathrm{Bool}}(G) \leq 4$. Finally, $\mathrm{ind}_2(G) = \dim_{\mathrm{geom}}(G) = \dim_{\mathrm{Bool}}(G) = 4$.

From Corollary 2.3 above, we deduce the following result.

Theorem 2.4. *The Boolean dimension of a path on n vertices ($n \in \mathbb{N}, n > 0$) is $n-1$. Every other n-vertex graph, with $n \geq 2$, has dimension at most $n-2$.*

Proof: Let P_n be the path on $\{0, \ldots, n-1\}$, whose edges are pairs $\{i, i+1\}$, with $i < n-1$. Suppose $n \geq 2$. Since P_n is the Boolean sum of its edges, $\dim_{\mathrm{Bool}}(P_n) \leq n-1$. Let $A := \{0, \ldots, n-2\}$. Then A is independent $\pmod 2$. Indeed, let X be a nonempty subset of A and x be its largest element, then the vertex $v := x+1$ is such that $|N_{P_n}(v) \cap X| = 1$. Thus $\mathrm{ind}_2(P_n) \geq n-1$. From the inequalities $n-1 \leq \mathrm{ind}_2(P_n) \leq \dim_{\mathrm{geom}}(P_n) \leq \dim_{\mathrm{Bool}}(P_n) \leq n-1$, the fact that the dimension of P_n is $n-1$ follows.

Now we prove that if the Boolean dimension of a graph G on n vertices is $n-1$, then G is a path. Observe first that G is connected. Otherwise, G is the direct sum $G' \oplus G''$ of two non trivial graphs G' and G'' with respectively n' and n'' vertices. As it is immediate to see, $\dim_{\text{Bool}}(G) = \dim_{\text{Bool}}(G' \oplus G'') \leq \dim_{\text{Bool}}(G') + \dim_{\text{Bool}}(G'') \leq n'-1+n''-1 = n-2$. Next we observe that G cannot be a cycle. Indeed, an easy induction shows that cycles on n vertices have dimension at most $n-2$. Indeed, the cycle C_3 is a clique thus has dimension 1. For $n \geq 4$, the cycle C_n on n vertices $\{0, \ldots, n-1\}$ is the Boolean sum of the cycle on the first $n-1$ vertices and the 3-vertex cycle on $\{0, n-2, n-1\}$, thus its dimension is at most $n-2$ (in fact it is equal to $n-2$; this is obvious for C_4 while for $n \geq 5$, its dimension is at least $n-2$ since it contains a path on $n-1$ vertices). Next, we check that if G has no more than four vertices, then it is a path. For the final step, we argue by induction, but we need a notation. Let $G := (V, E)$ be a graph and $x \in V$. Let G_{-x} be the subgraph of G induced by $V \setminus \{x\}$. Let $G^x := (G_{-x}) \dot{+} K_G(x)$, where $K_G(x) := K^{V \setminus \{x\}}_{N_G(x)}$. Let $\dot{G}^x$ be the graph obtained by adding to G^x the vertex x as an isolated vertex. In simpler terms, we obtain G^x by deleting from G the vertex x and by adding, via the Boolean sum, all edges between vertices of $N_G(x)$. For an example, if G is a path then G^x is a path on $V \setminus \{x\}$.

Claim 2.5. *If V is finite then* $|\dim_{\text{Bool}}(G) - \dim_{\text{Bool}}(G^x)| \leq 1$.

Proof of Claim 2.5 Note that $\dot{G}^x \dot{+} K^V_{N_G(x)\cup\{x\}} = G$ and $G \dot{+} K^V_{N_G(x)\cup\{x\}} = \dot{G}^x$. Thus $|\dim_{\text{Bool}}(G) - \dim_{\text{Bool}}(\dot{G}^x)| \leq \dim_{\text{Bool}}(K^V_{N_G(x)\cup\{x\}})$. Since $K_{N_G(x)\cup\{x\}}$ is a clique, its Boolean dimension is 1; and since $\dot{G}^x$ and G^x differ by an isolated vertex, they have the same Boolean dimension. The claimed inequality follows.

Now, let G be our graph on n vertices such that $\dim_{\text{Bool}}(G) = n-1$. Suppose that every graph G' on n' vertices, $n' < n$, is a path whenever $\dim_{\text{Bool}}(G') = n'-1$.

Claim 2.6. *G^x is a path for every $x \in V(G)$.*

Indeed, since G^x has $n-1$ vertices, $\dim_{\text{Bool}}(G^x) \leq n-2$; since $\dim_{\text{Bool}}(G) = n-1$, the claim above ensures that $\dim_{\text{Bool}}(G^x) = n-2$. The conclusion follows for the hypothesis on graphs with $n-1$ vertices.

Claim 2.7. *Let $x, y \in V(G)$ with $x \neq y$. If G^x and G^y are two paths P_x and P_y, then $d_G(x), d_G(y) \leq 2$ if $\{x, y\} \notin E(G)$, and $d_G(x), d_G(y) \leq 3$ otherwise.*

Proof of Claim 2.7 We have $G^x \dot{+} G^y = G \dot{+} K^V_{N_G(x)\cup\{x\}} \dot{+} G \dot{+} K^V_{N_G(y)\cup\{y\}} = K^V_{N_G(x)\cup\{x\}} \dot{+} K^V_{N_G(y)\cup\{y\}}$. Since G^x and G^y are two paths P_x and P_y, $P_x \dot{+} P_y = K^V_{N_G(x)\cup\{x\}} \dot{+} K^V_{N_G(y)\cup\{y\}}$. We have $d_{P_x \dot{+} P_y}(x) \leq 2$, hence for the Boolean sum $E_G(x) \dot{+} E_G(y)$ of stars $E_G(x)$ and $E_G(y)$, we have $d_{E_G(x) \dot{+} E_G(y)}(x) \leq 2$. The conclusion of the claim follows.

Now let $x \in V$. Since $d_G(x) \leq 3$ and $n \geq 5$ there is some vertex y not linked to x by an edge. Hence by Claim 2.7, $d_G(x) \leq 2$. From this follows that G is a direct sum of paths and cycles.

Since G must be connected and cannot be a cycle, G is a path. □

We thank Bondy (2009) for suggesting this result several years ago. In fact, it is a consequence of previous results about geometric dimension of graphs, obtained for general fields (Bento and Leal Duarte, 2005; Rheinboldt and Shepherd, 1974).

We go from paths to trees as follows.

Definition 2.8. Let $T := (V, E)$ be a tree. A *star decomposition* Σ of T is a family $\{S_1, \ldots, S_k\}$ of subtrees of T such that each S_i is isomorphic to $K_{1,m}$ (a star) for some $m \geq 1$, the stars are mutually edge-disjoint, and each edge of T is an edge of some S_i. For a star decomposition Σ, let $t(\Sigma)$ be the number of trivial stars in Σ (stars that are isomorphic to $K_{1,1}$), and let $s(\Sigma)$ be the number of nontrivial stars in Σ (stars that are isomorphic to $K_{1,m}$ for some $m > 1$). We define the parameter $m(T) \coloneqq \min_{\Sigma}\{t(\Sigma) + 2s(\Sigma)\}$ over all star decompositions Σ of T. A star decomposition Σ of T for which $t(\Sigma) + 2s(\Sigma) = m(T)$ is called an *optimal star decomposition* of T.

The Boolean dimension of a graph counts the minimum number of cliques needed to obtain this graph as a Boolean sum. If $\Sigma \coloneqq \{S_1, \dots, S_k\}$ is a star decomposition of a tree T, one has $\dim_{\text{Bool}}(T) \leq \sum_{i=1}^{k} \dim_{\text{Bool}}(S_i)$. Since $\dim_{\text{Bool}}(S_i) = 1$ if S_i is a trivial star, and $\dim_{\text{Bool}}(S_i) = 2$ otherwise (note that if $S_i = K_{1,m}$, it is the Boolean sum of a clique on $m+1$ vertices and a clique on a subset of m vertices), hence we have $\sum_{i=1}^{k} \dim_{\text{Bool}}(S_i) = t(\Sigma) + 2s(\Sigma)$, hence $\dim_{\text{Bool}}(T) \leq t(\Sigma) + 2s(\Sigma)$. The inequality $\dim_{\text{Bool}}(T) \leq m(T)$ follows.

Here is our result.

Theorem 2.9. *For all trees T, we have* $\text{ind}_2(T) = \dim_{\text{Bool}}(T) = m(T)$.

We introduce the following definition.

Definition 2.10. A *cherry* in a tree T is a maximal subtree S isomorphic to $K_{1,m}$ for some $m > 1$ that contains m end vertices of T. We refer to a cherry with m edges as an m-cherry.

Proposition 2.11. *Let $T \coloneqq (V, E)$ be a tree that contains a cherry. If all proper subtrees T' of T satisfy* $\text{ind}_2(T') = m(T')$*, then* $\text{ind}_2(T) = m(T)$.

Proof: Let $x \in V$ be the center of a k-cherry in T, with $N_T(x) = \{u_1, \dots, u_k, w_1, \dots, w_\ell\}$, where $d_T(u_i) = 1$ for all i, and $d_T(w_i) > 1$ for all i. For each $i = 1$ to ℓ, let T_i be the maximal subtree that contains w_i but does not contain x.

First, we show that any optimal star decomposition of T in which x is not the center of a nontrivial star can be transformed into an optimal star decomposition in which x is the center of a nontrivial star. Consider an optimal star decomposition Σ in which x is not the center of a nontrivial star. Therefore, edges xu_i are trivial stars of Σ. Now if $k > 2$ or if there is a trivial star xw_i in Σ, then we could have improved $t(\Sigma) + 2s(\Sigma)$ by replacing all trivial stars containing x by their union, which is a star centered at x. Hence, assume that $k = 2$ and each w_i is the center of a nontrivial star S_i, which contains the edge xw_i. Now replace each S_i by $S'_i \coloneqq S_i - xw_i$, and add a new star centered at x with edge set $\{xw_1, \dots, xw_\ell, xu_1, xu_2\}$. The new decomposition is also optimal.

Now consider an optimal star decomposition Σ in which x is the center of a nontrivial star. The induced decompositions on T_i are all optimal since Σ is optimal. For each $i \in \{1, \dots, \ell\}$, let A_i be a maximum size independent set in T_i. Hence $|A_i| = \text{ind}_2(T_i) = m(T_i)$ for all $i \geq 1$, and $m(T) = 2 + \sum_i m(T_i) = 2 + \sum_i \text{ind}_2(T_i)$. We show that $A \coloneqq \{x, u_1\} \cup (\cup_i A_i)$ is a maximum size independent set in T.

Consider a non-empty set $X \subseteq A$. We show that there exists $v \in V \setminus X$ such that $|N_T(v) \cap X|$ is odd. If $x \in X$, we have $N_T(u_2) \cap X = \{x\}$. If $X = \{u_1\}$, we have $N_T(x) \cap X = \{u_1\}$. So suppose $x \notin X$ and $X \neq \{u_1\}$. Let $B_i \coloneqq X \cap V(T_i)$ for $i \in \{1, \dots, \ell\}$. Since B_i is nonempty for some i, and $x \notin X$, we find $v \in V(T_i) \setminus B_i$ such that $|N_{T_i}(v) \cap B_i|$ is odd. Now $|N_T(v) \cap X|$ is odd since $x \notin X$ and v is not adjacent to u_1. Moreover, $|A| = m(T)$. □

Proposition 2.12. *Let $T \coloneqq (V, E)$ be a tree that contains a vertex y of degree 2 adjacent to a vertex z of degree 1. If* $\text{ind}_2(T - z) = m(T - z)$*, then* $\text{ind}_2(T) = m(T)$.

Proof: First, we show that $m(T) = m(T - z) + 1$. If there is an optimal star decomposition of $T - z - y$ in which some vertex x is the center of a star, then $m(T - z) = m(T - z - y)$ and $m(T) = m(T - z) + 1$, else $m(T - z) = m(T - z - y) + 1$ and $m(T) = m(T - z - y) + 2$.

Now we consider a maximum sized independent set A' in $T-z$. We have $|A'| = \mathrm{ind}_2(T-z) = m(T-z)$. We define $A \coloneqq A' \cup \{y\}$ if $y \notin A'$; and $A \coloneqq A' \cup \{z\}$ if $y \in A'$. We show that A is independent in T.

Case 1: $y \notin A'$, hence $y \in A$ and $z \notin A$. Let $B \subseteq A, B \neq \varnothing$.

If $y \in B$, then $|N_T(z) \cap B|$ is odd.

If $y \notin B$, then $B \subseteq A'$, hence there exists $v \in V(T-z)$ such that $|N_{T-z}(v) \cap B|$ is odd, and $|N_T(v) \cap B|$ is odd.

Case 2: $y \in A'$, hence $z \in A$. Let $B \subseteq A, B \neq \varnothing$.

If $z \notin B$, then $B \subseteq A'$. Find $v \in V(T - z) \smallsetminus B$ such that $|N_{T-z}(v) \cap B|$ is odd. Hence $|N_T(v) \cap B|$ is odd.

Now suppose that $z \in B$. If $B = \{z\}$, then $|N_T(y) \cap B|$ is odd. Otherwise, consider $B \smallsetminus \{z\}$, which is a subset of A'. Find $v \in V(T - z) \smallsetminus (B \smallsetminus \{z\})$ such that $|N_{T-z}(v) \cap (B \smallsetminus \{z\})|$ is odd. If $v \neq y$, then $|N_T(v) \cap B|$ is odd. If $v = y$ and $x \in N_T(y) \smallsetminus \{z\}$, then $|N_T(v) \cap B|$ is even and $x \in B$. In this case, let $B' \coloneqq (B \smallsetminus \{z\}) \cup \{y\}$. This is a subset of A'. Find $u \in V(T - z) \smallsetminus B'$ such that $|N_{T-z}(u) \cap B'|$ is odd. Since B' contains x and y, we conclude that u is not adjacent to any of y and z, hence $|N_T(u) \cap B|$ is odd.

Thus we have shown that A is independent. We have $\mathrm{ind}_2(T) \geq |A| = |A'| + 1 = m(T - z) + 1 = m(T)$. Since $\mathrm{ind}_2(T)$ cannot be more than $m(T)$, we have $\mathrm{ind}_2(T) = m(T)$. □

Proof Proof of Theorem 2.9: If a tree T has two vertices, then $\mathrm{ind}_2(T) = m(T) = 1$. Each tree with at least 3 vertices contains a cherry or a vertex of degree 2 adjacent to a vertex of degree 1. (This is seen by considering the second-to-last vertex of a longest path in T.) Now, induction on the number of vertices, using Propositions 2.11 and 2.12, implies the result. □

3 Inversion index of a tournament and Boolean dimension

3.1 Inversion index of a tournament

Let T be a tournament. Let $V(T)$ be its vertex set and $A(T)$ be its arc set. An *inversion* of an arc $a \coloneqq (x, y) \in A(T)$ consists to replace the arc a by $a^\star \coloneqq (y, x)$ in $A(T)$. For a subset $X \subseteq V(T)$, let $Inv(T, X)$ be the tournament obtained from T after reversing all arcs $(x, y) \in A(T) \cap (X \times X)$. For example, $Inv(T, V)$ is T^*, the *dual* of T. For a finite sequence $(X_i)_{i<m}$ of subsets of $V(T)$, let $Inv(T, (X_i)_{i<m})$ be the tournament obtained from T by reversing successively all the arcs in each of the subsets X_i, $i < m$, that is, the tournament equal to T if $m = 0$ and to $Inv(Inv(T, (X_i)_{i<m-1}), X_{m-1})$ if $m \geq 1$. Said differently, an arc $(x, y) \in A(T)$ is reversed if and only if the number of indices i such that $\{x, y\} \subseteq X_i$ is odd. The *inversion index* of T, denoted by $i(T)$, is the least integer m such that there is a sequence $(X_i)_{i<m}$ of subsets of $V(T)$ for which $Inv(T, (X_i)_{i<m})$ is acyclic.

In the sequel, we consider tournaments for which this index is finite. In full generality, the inversion index of a tournament T can be defined as the least cardinal κ such the Boolean sum of T and a graph of Boolean dimension κ is acyclic. The case κ finite is stated in Lemma 3.8 below. We leave tournaments with infinite inversion index to further studies.

The motivation for the notion of inversion index originates in the study of critical tournaments. Indeed, the critical tournaments of Schmerl and Trotter (1993) can be easily defined from acyclic tournaments by means of one or two inversions whereas the (-1)-critical tournaments, characterized in Belkhechine et al. (2007), can be defined by means of two, three or four inversions Belkhechine (2009). Another interest comes from the point of view of logic.

Results about the inversion index originate in the thesis of Belkhechine (2009). Some results have been announced in Belkhechine et al. (2010); they have been presented at several conferences by the first author and included in a circulating manuscript Belkhechine et al. (2012). The lack of answer for some basic questions is responsible for the delay of publication.

The inversion index is a variant of the *Slater index*: the least number of arcs of a tournament which have to be reversed in order to get an acyclic tournament (Slater, 1961). The complexity of the computation of the Slater index was raised by Bang-Jensen and Thomassen (1992). Alon (2006) and independently Charbit, Thomassé, and Yeo (2007) showed that the problem is NP-hard. An extension of the inversion index to oriented graphs is studied in Bang-Jensen et al. (2020).

Problem 3.1. Is the computation of the inversion index NP-hard?

Question 3.2. Are there tournaments of arbitrarily large inversion index?

This last question has a positive answer. There are two reasons, the first one is counting, the second one, easier, is based on the notion of well-quasi-ordering.

For $n \in \mathbb{N}$, let $i(n)$ be the maximum of the inversion index of tournaments on n vertices. We have $i(n) = 0$ for $n \leq 2$, $i(3) = i(4) = 1$, $i(5) = i(6) = 2$. For larger n a counting argument Belkhechine (2009); Belkhechine et al. (2010, 2012) yields the following result.

Theorem 3.1. $\frac{n-1}{2} - \log_2 n \leq i(n) \leq n - 4$ *for all integer* $n \geq 6$.

It is quite possible that $i(n) \geq \lfloor \frac{n-1}{2} \rfloor$, due to the path of strong connectivity (it is not even known if reverse inequality holds).

The *path of strong connectivity on* n vertices is the tournament T_n defined on $\mathbb{N}_{<n} := \{0, \ldots, n-1\}$ whose arcs are all pairs $(i, i+1)$ and (j, i) such that $i + 1 < n$ and $i < j < n$.

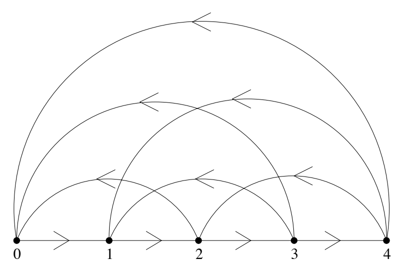

Fig. 1: Path of strong connectivity on 5 vertices

Question 3.3. Is the inversion index of a path of strong connectivity on n vertices equal to $\lfloor \frac{n-1}{2} \rfloor$?

3.2 Well-quasi-ordering

Basic notions of the theory of relations apply to the study of the inversion index. These notions include the quasi order of embeddability, the hereditary classes and their bounds, and the notion of well-quasi-order. For those, we refer to Fraïssé (2000).

Let $\mathcal{I}_m^{<\omega}$ be the class of finite tournaments T whose inversion index is at most m. This is a hereditary class in the sense that if $T \in \mathcal{I}_m^{<\omega}$ and T' is embeddable into T then $T' \in \mathcal{I}_m^{<\omega}$. It can be characterized by obstructions or bounds. A *bound* is a tournament not in $\mathcal{I}_m^{<\omega}$ such that all proper subtournaments are in $\mathcal{I}_m^{<\omega}$. We may note that the inversion index of every bound of $\mathcal{I}_m^{<\omega}$ is at least $m+1$. Hence, the fact that $\mathcal{I}_m^{<\omega}$ is distinct of the class of all finite tournaments provides tournaments of inversion index larger than m. This fact relies on the notion of well-quasi-ordering.

A poset P is *well-quasi-ordered* if every sequence of elements of P contains an increasing subsequence.

Theorem 3.2. *The class of all finite tournaments is not well-quasi-ordered by embeddability.*

This is a well known fact. As indicated by a referee, it has been mentioned by several authors. See e.g., Latka (1994) for a much stronger version of Theorem 3.2 and also subsection 3.1 of Cherlin and Latka (2000). For the convenience of the reader we give a proof.

Proof: Let T_n be the path of strong connectivity on $\{0,\dots,n-1\}$ as defined above. Let C_n be the tournament obtained from T_n by reversing the arc $(n-1,0)$. We claim that for $n \geq 7$, the C_n's form an antichain. Indeed, to C_n we may associate the 3-uniform hypergraph H_n on $\{0,\dots,n-1\}$ whose 3-element hyperedges are the 3-element cycles of C_n. An embedding from some C_n to another C_m, $m > n$, induces an embedding from H_n to H_m. To see that such an embedding cannot exist, observe first that the vertices 0 and $n-1$ belong to exactly $n-3$ hyperedges, and the vertices 1 and $n-2$ belong to exactly two hyperedges, the other vertices to three hyperedges, hence an embedding h will send $\{0,n-1\}$ on $\{0,m-1\}$. The preservation of the arc $(0,n-1)$ imposes $h(0)=0$ and $h(n-1)=m-1$. Then, the preservation of the arcs $(i,i+1)$ yields a contradiction since $n<m$. □

Theorem 3.3. *Belkhechine et al. (2010) For each $m \in \mathbb{N}$, the class $\mathcal{I}_m^{<\omega}$ is well-quasi-ordered.*

Proof: The class $\mathcal{L}_m^{<\omega}$ made of a finite linear order L with m unary predicates $U_1,\dots,U_m$ (alias m distinguished subsets) and ordered by embeddability is well-quasi-ordered. This is a straightforward consequence of Higman's theorem on words (see Higman, 1952), in fact, an equivalent statement. Higman's result asserts that the collection of words on a finite alphabet, ordered by the subword ordering, is well-quasi-ordered. Since members of $\mathcal{L}_m^{<\omega}$ can be coded by words on an alphabet with 2^m elements, the class $\mathcal{L}_m^{<\omega}$ is well-quasi-ordered. The map associating to each $(L,U_1,\dots,U_m)$ the Boolean sum $L\dot{+}U_1\dots\dot{+}U_m$ preserves the embeddability relation, hence the range of that map is well-quasi-ordered. This range being equal to $\mathcal{I}_m^{<\omega}$, this later class is well-quasi-ordered. □

Corollary 3.4. *There are finite tournaments with arbitrarily large inversion index.*

We have the following result concerning the bounds.

Theorem 3.5. *Belkhechine et al. (2010) The class $\mathcal{I}_m^{<\omega}$ has only finitely many bounds.*

Proof: From the proof of Theorem 3.3, the class $\mathcal{I}_{m,1}^{<\omega}$ made of tournaments of $\mathcal{I}_m^{<\omega}$, with one unary predicate added, is well-quasi-ordered. According to an adaptation of Proposition 2.2 of Pouzet (1972) translated in this case, $\mathcal{I}_m^{<\omega}$ has finitely many bounds. □

We thank the referee for observing that the well-quasi-ordering of $\mathcal{I}^{<\omega}_{m,1}$ suffices to yield the finiteness of the bounds of $\mathcal{I}^{<\omega}_{m}$.

Question 3.4. What is the maximum of the cardinality of bounds of $\mathcal{I}^{<\omega}_{m}$?

Remark 3.6. It must be observed that the collection of graphs with geometric dimension at most m over a fixed finite field has finitely many bounds and an upper bound on their cardinality is given in Ding and Kotlov (2006). How the cardinality of these bounds relate to the cardinality of bounds of $\mathcal{I}^{<\omega}_{m}$ is not known.

3.3 Boolean dimension and concrete examples of tournaments with large inversion index

Let $C_3.\underline{n}$ be the sum of copies of the 3-cycle C_3 indexed by the n-element acyclic tournament $\underline{n} := (\{0, \dots, n-1\}, \{(i,j) \mid 0 \leq i < j \leq n-1\})$ with $0 < \dots < n-1$.

Theorem 3.7. *The inversion index of the sum $C_3.\underline{n}$ of 3-cycles over an n-element acyclic tournament is n.*

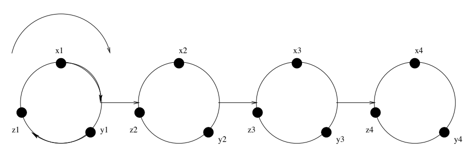

Fig. 2: $C_3.\underline{4}$

No elementary proof is known. The proof we present relies on the notion of Boolean sum of graphs. According to the definition of Boolean sum, we have the following result immediately.

Lemma 3.8. *The inversion index of a tournament T is equal to the least integer k such that the Boolean sum $T\dot{+}G$ of T with a graph G of Boolean dimension k is an acyclic tournament.*

Proof of Theorem 3.7. Let $T := C_3.\underline{n}$, $V := V(T)$ and $r := i(T)$. Clearly $r \leq n$. Conversely, let H be a graph with vertex set V such that $L := T\dot{+}H$ is an acyclic tournament and $\dim_{\mathrm{Bool}}(H) = r$. Let $U := (\mathbb{F}_2)^r$ equipped with the ordinary scalar product $\mid$ and $f : V \to U$ be a representation of H.

Claim 3.9. *For each $i \in \{0, \dots, n-1\}$, we may enumerate the vertices of $\{0,1,2\} \times \{i\}$ into x_i, y_i, z_i in such a way that $(x_i, y_i), (y_i, z_i), (z_i, x_i)$ are arcs of T, $(f(x_i)|f(z_i)) = 1$ and $(f(x_i)|f(y_i)) = 0$.*

Claim 3.10. *The set $\{f(x_i) : i < n\}$ is linearly independent in U.*

Proof of Claim 3.10. This amounts to prove that $\sum_{i \in I} f(x_i) \neq 0$ for every non-empty subset I of $\{0, \dots, n-1\}$. Let I be such a subset. Let $m \in \{0, \dots, n-1\}$ such that x_m is the largest element of $\{x_i : i \in I\}$ in the acyclic tournament L.

Subclaim 3.11. $(f(x_i)|f(z_m)) = (f(x_i)|f(y_m))$ for each $i \in I \setminus \{m\}$.

Proof of Subclaim 3.11. By construction, we have $x_m <_L z_m$ and $x_m <_L y_m$, hence by transitivity $x_i <_L z_m$ and $x_i <_L y_m$. If $i < m$ in the natural order then, by definition of T, $(x_i, z_m) \in A(T)$ and $(x_i, y_m) \in A(T)$, thus $(f(x_i)|f(z_m)) = 0 = (f(x_i)|f(y_m))$, whereas if $i > m$ in the natural order, then

$(z_m, x_i) \in A(T)$ and $(y_m, x_i) \in A(T)$, thus $(f(x_i)|f(z_m)) = 1 = (f(x_i)|f(y_m))$, proving the subclaim. □

Since $(f(x_m)|f(z_m)) = 1$ and $(f(x_m)|f(y_m)) = 0$, it follows that $\sum_{i \in I}(f(x_i)|f(z_m)) \neq \sum_{i \in I}(f(x_i)|f(y_m))$. That is $((\sum_{i \in I} f(x_i))|f(z_m)) \neq ((\sum_{i \in I} f(x_i))|f(y_m))$. Thus the sum $\sum_{i \in I} f(x_i) \neq 0$ as claimed. □

We have $n \leq r$. This proves the theorem. □

Acknowledgements

The authors thank a referee of this paper for his extremely careful and detailed suggestions which significantly improved the presentation of this paper and also the content (e.g., Theorem 1.17). The third author would like to thank Institut Camille Jordan, Université Claude Bernard Lyon 1 for hospitality and support. Support from CAPES Brazil (Processo: : 88887.364676/2019-00) and Labex MILYON: ANR-10-LABX-0070 are gratefully acknowledged. The first author is pleased to thank Uri Avraham, Adrian Bondy and Christian Delhommé for their contribution.

References

N. Alon. Ranking tournaments. *SIAM J. Discrete Math.*, 20:137–142, 2006.

U. Avraham. Personal communication, April 2021.

J. Bang-Jensen and C. Thomassen. A polynomial algorithm for the 2-path problem for semicomplete digraphs. *SIAM J. Discrete Math.*, 5(3):366–376, 1992.

J. Bang-Jensen, J. C. Ferreira da Silva, and F. Havet. Inversion number of an oriented graph and related parameters. In M. Couceiro, P. Monnin, and A. Napoli, editors, *Proceedings of the 1st International Conference on Algebras, Graphs and Ordered Sets (ALGOS 2020)*, pages 119–139, Nancy, France, Aug. 2020. HAL. URL `https://hal.inria.fr/hal-02918958`.

H. Belkhechine. *Indécomposabilité des graphes et des tournois*. PhD thesis, Université de Sfax and Université Claude-Bernard Lyon 1, Jul 2009.

H. Belkhechine, I. Boudabbous, and J. Dammak. Morphologie des tournois (-1)-critiques. *C. R. Acad. Sci. Paris, Ser. I*, 345:663–666, 2007.

H. Belkhechine, M. Bouaziz, I. Boudabbous, and M. Pouzet. Inversion dans les tournois. *C.R. Acad. Sci. Paris, Ser. I*, 348:703–707, 2010.

H. Belkhechine, M. Bouaziz, I. Boudabbous, and M. Pouzet. Inversions in tournaments. Unpublished manuscript, 2012.

A. Bento and A. Leal Duarte. On Fiedler's characterization of tridiagonal matrices over arbitrary fields. *Linear Algebra Appl.*, 401:467–481, 2005.

A. Bondy. Personal communication, July 2009.

P. Charbit, S. Thomassé, and A. Yeo. The minimum feedback arc set problem is NP-hard for tournaments. *Combin. Probab. Comput.*, 16(1):1–4, 2007.

G. Cherlin and B. Latka. Minimal antichains in well-founded quasi-orders with an application to tournaments. *J. Combin. Theory Ser. B*, 80(2):258–276, 2000.

P. M. Cohn. *Algebra*, volume 3. John Wiley, Sons, Ltd., Chichester, second edition, 1991. xii+474 pp.

B. Courcelle and C. Delhommé. The modular decomposition of countable graphs. Definition and construction in monadic second-order logic. *Theoretical Computer Science*, 394(2):1–38, 2008.

C. Delhommé. Personal communication, April 2021.

G. Ding and A. Kotlov. On minimal rank over finite fields. (English summary). *Electron. J. Linear Algebra*, 15:210–214, 2006.

S. M. Fallat and L. Hogben. The minimum rank of symmetric matrices described by a graph: a survey. *Linear Algebra and its Applications*, 426(2–3):558–582, 2007.

R. Fraïssé. Theory of relations. Studies in Logic and the Foundations of Mathematics, 145. North-Holland Publishing Co., Amsterdam, revised edition, 2000. xii+451 pp.

M. H. Garzon. Symplectic embeddings of graphs. *Journal of Combinatorial Mathematics and Combinatorial Computing*, 2:193–207, 1987.

C. D. Godsil and G. F. Royle. Chromatic number and the 2-rank of a graph. *Journal of Combinatorial Theory, Series B*, 81(1):142–149, 2001a.

C. D. Godsil and G. F. Royle. *Algebraic graph theory*. Graduate Texts in Mathematics, 207. Springer-Verlag, New York, 2001b. xx+439 pp.

J. Grout. The minimum rank problem over finite fields. *Electron. J. Linear Algebra*, 20:691–716, 2010.

G. Higman. Ordering by divisibility in abstract algebras. *Proc. London Math. Soc.*, 3(2):326–336, 1952.

T. Jech. *Set theory. The third millennium edition, revised and expanded*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. xiv+769 pp.

K. Kunen. *Set theory. Studies in Logic, 34*. College Publications. London, 2011. viii+401 pp.

B. Latka. Finitely constrained classes of homogeneous directed graphs. *J. Symbolic Logic*, 59(1):124–139, 1994.

M. Pouzet. Un bel ordre d'abritement et ses rapports avec les bornes d'une multirelation. *C. R. Acad. Sci. Paris Sér. A-B*, 274:A1677–A1680, 1972.

W. Rheinboldt and R. Shepherd. On a characterization of tridiagonal matrices by M. Fiedler. *Linear Algebra Appl.*, 8:87–90, 1974.

A. Rinot. The Δ-system lemma: an elementary proof. URL `https://blog.assafrinot.com`.

J. H. Schmerl and W. T. Trotter. Critically indecomposable partially ordered sets, graphs, tournaments and other binary relational structures. *Discrete Math.*, 113:191–205, 1993.

P. Slater. Inconsistencies in a schedule of paired comparison. *Biomathematica*, 48:303–312, 1961.

J. H. van Lint and R. M. Wilson. *A Course in Combinatorics*. Cambridge University Press, second edition, 2001. xiv+602 pp.