Shattering-extremal set systems of VC dimension at most 2

Tamás Mészáros

Department of Mathematics, Central European University
Institute of Mathematics, Budapest University of Technology and Economics
tmeszaros87@gmail.com

Lajos Rónyai[1]

Computer and Automation Research Institute, Hungarian Academy of Sciences
Institute of Mathematics, Budapest University of Technology and Economics
lajos@ilab.sztaki.hu

Abstract

We say that a set system $\mathcal{F} \subseteq 2^{[n]}$ shatters a given set $S \subseteq [n]$ if $2^S = \{F \cap S : F \in \mathcal{F}\}$. The Sauer inequality states that in general, a set system $\mathcal{F}$ shatters at least $|\mathcal{F}|$ sets. Here we concentrate on the case of equality. A set system is called shattering-extremal if it shatters exactly $|\mathcal{F}|$ sets. In this paper we characterize shattering-extremal set systems of Vapnik-Chervonenkis dimension 2 in terms of their inclusion graphs, and as a corollary we answer an open question from [11] about leaving out elements from shattering-extremal set systems in the case of families of Vapnik-Chervonenkis dimension 2.

1 Introduction

Throughout this paper n will be a positive integer, the set $\{1, 2, \ldots, n\}$ will be referred to shortly as $[n]$ and the power set of any set $S \subseteq [n]$ will be denoted by 2^S. For a set system $\mathcal{F} \subseteq 2^{[n]}$ we will write $supp(\mathcal{F})$ for its support, i.e. $supp(\mathcal{F}) = \bigcup_{F \in \mathcal{F}} F$.

The central notion of our study is *shattering*.

Definition 1.1 *A set system $\mathcal{F} \subseteq 2^{[n]}$ shatters a given set $S \subseteq [n]$ if*

$$2^S = \{F \cap S : F \in \mathcal{F}\}.$$

The family of subsets of $[n]$ shattered by $\mathcal{F}$ is denoted by $Sh(\mathcal{F})$. The following inequality states that in general, a set system $\mathcal{F}$ shatters at least $|\mathcal{F}|$ sets.

Proposition 1.1 $|Sh(\mathcal{F})| \geq |\mathcal{F}|$ *for every set system* $\mathcal{F} \subseteq 2^{[n]}$.

The statement was proved by several authors, e.g. Aharoni and Holzman [1], Pajor [12], Sauer [13], Shelah [14]. Often it is referred to as the Sauer inequality. Here we are interested in the case of equality.

Definition 1.2 *A set systems* $\mathcal{F} \subseteq 2^{[n]}$ *is shattering-extremal, or s-extremal for short, if it shatters exactly* $|\mathcal{F}|$ *sets, i.e.* $|\mathcal{F}| = |Sh(\mathcal{F})|$.

Many interesting results have been obtained in connection with these combinatorial objects, among others by Bollobás, Leader and Radcliffe in [3], by Bollobás and Radcliffe in [4], by Frankl in [5]. Füredi and Quinn in [6], and recently Kozma and Moran in [8] provided interesting examples of s-extremal set systems. Anstee, Rónyai and Sali in [2] related shattering to standard monomials of vanishing ideals, and based on this relation, the present authors in [9] and in [10] developed algebraic methods for the investigation of s-extremal families.

Definition 1.3 *The inclusion graph of a set system* $\mathcal{F} \subseteq 2^{[n]}$, *denoted by* $\mathbb{G}_{\mathcal{F}}$, *is the simple directed edge labelled graph whose vertices are the elements of* $\mathcal{F}$, *and there is a directed edge with label* $j \in [n]$ *going from* G *to* F *exactly when* $F = G \cup \{j\}$.

$\mathbb{G}_{\mathcal{F}}$ is actually the Hasse diagram of the poset $\mathcal{F}$ with edges directed and labelled in a natural way. The inclusion graph of the complete set system $2^{[n]}$ will be denoted by $\mathbb{H}_n$. The undirected version of $\mathbb{H}_n$ is often referred to as the Hamming graph $H(n, 2)$, or as the hypercube of dimension n, whose vertices are all $0 - 1$ vectors of length n, and two vertices are adjacent iff they differ in exactly one coordinate. When computing distances between vertices in the inclusion graph $\mathbb{G}_{\mathcal{F}}$ we forget about the direction of edges, and define the distance between vertices $F, G \in \mathcal{F}$, denoted by $d_{\mathbb{G}_{\mathcal{F}}}(F, G)$, as their graph distance in the undirected version of $\mathbb{G}_{\mathcal{F}}$, i.e. the length of the shortest path between them in the undirected version of $\mathbb{G}_{\mathcal{F}}$. Similarly, some edges in $\mathbb{G}_{\mathcal{F}}$ form a path between two vertices if they do so in the undirected version of $\mathbb{G}_{\mathcal{F}}$. For example, the distance between two vertices $F, G \subseteq [n]$ in $\mathbb{H}_n$ is just

the size of the symmetric difference $F \triangle G$, i.e. $d_{\mathbb{H}_n}(F, G) = |F \triangle G|$. As a consequence, when only distances of vertices will be considered, and the context will allow, we omit the directions of edges to avoid unnecessary case analysis, and will specify edges by merely listing their endpoints.

Definition 1.4 *The Vapnik-Chervonenkis dimension of a set system $\mathcal{F} \subseteq 2^{[n]}$, denoted by $dim_{VC}(\mathcal{F})$, is the maximum cardinality of a set shattered by $\mathcal{F}$.*

The general task of giving a good description of s-extremal systems seems to be too complex at this point, therefore we restrict our attention to the simplest cases, where the VC-dimension of $\mathcal{F}$ is small. S-extremal systems, where the VC-dimension is at most 1 were fully described in [11].

Proposition 1.2 *(See [11].) A set system $\mathcal{F} \subseteq 2^{[n]}$ is s-extremal and of VC dimension at most 1 iff $G_{\mathcal{F}}$ is a tree and all labels on the edges are different.*

Proposition 1.2 can also be interpreted as follows:

Proposition 1.3 *(See [11]) There is a one-to-one correspondence between s-extremal families $\mathcal{F} \subseteq 2^{[n]}$ of VC-dimension 1 with $supp(\mathcal{F}) = [n]$, $\cap_{F \in \mathcal{F}} F = \emptyset$ and directed edge-labelled trees on $n+1$ vertices, all edges having a different label from $[n]$.*

Note that the assumptions $supp(\mathcal{F}) = [n]$ and $\cap_{F \in \mathcal{F}} F = \emptyset$ are not restrictive. Both of them can be assumed to hold without loss of generality, otherwise one could omit common elements and then restrict the ground set to $supp(\mathcal{F})$.

In this paper we continue the work initiated in [11], and characterize s-extremal set systems of VC-dimension at most 2. We do this by providing an algorithmic procedure for constructing the inclusion graphs of all such set systems. This characterization then allows us to answer an open question, posed in [11], about leaving out elements from such set systems.

The paper is organized as follows. After the introduction in Section 2 we investigate the properties of shattering and its connection to inclusion graphs. Next, in Section 3 we propose a building process for extremal families and investigate its properties. Based on this building process in Section 4 we present and prove our main results. Finally in Section 5 we make some concluding remarks concerning future work.

2   Preliminaries

To start with, we first introduce a useful subdivision of set systems.

Definition 2.1 *The standard subdivision of a set system $\mathcal{F} \subseteq 2^{[n]}$ with respect to an element $i \in [n]$ consists of the following two set systems:*

$$\mathcal{F}_0 = \{F : F \in \mathcal{F}; i \notin F\} \subseteq 2^{[n]\setminus\{i\}} \ and$$
$$\mathcal{F}_1 = \{F\setminus\{i\} : F \in \mathcal{F}; i \in F\} \subseteq 2^{[n]\setminus\{i\}}.$$

For the sake of completeness we provide a possible proof of Proposition 1.1, whose main idea will be useful later on.

Proof:(of Proposition 1.1) We will prove this statement by induction on n. For $n = 1$ the statement is trivial. Now suppose that $n > 1$, and consider the standard subdivision of $\mathcal{F}$ with respect to the element n. Note that $\mathcal{F}_0, \mathcal{F}_1 \subseteq 2^{[n-1]}$ and hence by the induction hypothesis we have $|Sh(\mathcal{F}_0)| \geq |\mathcal{F}_0|$ and $|Sh(\mathcal{F}_1)| \geq |\mathcal{F}_1|$. Moreover $|\mathcal{F}| = |\mathcal{F}_0| + |\mathcal{F}_1|$, $Sh(\mathcal{F}_0) \cup Sh(\mathcal{F}_1) \subseteq Sh(\mathcal{F})$ and if $S \in Sh(\mathcal{F}_0) \cap Sh(\mathcal{F}_1)$, then according to the definition of $\mathcal{F}_0$ and $\mathcal{F}_1$ we have $S \cup \{n\} \in Sh(\mathcal{F})$. Summarizing

$$|Sh(\mathcal{F})| \geq |Sh(\mathcal{F}_0)| + |Sh(\mathcal{F}_1)| \geq |\mathcal{F}_0| + |\mathcal{F}_1| = |\mathcal{F}|. \blacksquare$$

From the proof of Proposition 1.1 it is easy to see, that if $\mathcal{F}$ is s-extremal, then so are the systems $\mathcal{F}_0$ and $\mathcal{F}_1$ in the standard subdivision with respect to any element $i \in [n]$. Iterating this for an s-extremal system $\mathcal{F} \subseteq 2^{[n]}$ we get that for all pairs of sets $A \subseteq B \subseteq [n]$, the system

$$\{F\setminus A \mid F \in \mathcal{F}, \ A \subseteq F \subseteq B\}$$

is s-extremal. Moreover if in the above system we add A to every set, then the family of shattered sets remains unchanged, hence we get that the subsystem

$$\mathcal{F}_{A,B} = \{F \mid F \in \mathcal{F}, \ A \subseteq F \subseteq B\} \subseteq \mathcal{F}$$

is also s-extremal.

In [3] and [4] a different version of shattering, *strong shattering* is introduced .

Definition 2.2 *A set system $\mathcal{F} \subseteq 2^{[n]}$ strongly shatters the set $F \subseteq [n]$, if there exists $I \subseteq [n] \backslash F$ such that*

$$2^F + I = \{H \cup I \mid H \subseteq F\} \subseteq \mathcal{F}.$$

The family of all sets strongly shattered by some set system $\mathcal{F}$ is denoted by $st(\mathcal{F})$. Clearly $st(\mathcal{F}) \subseteq Sh(\mathcal{F})$, both $Sh(\mathcal{F})$ and $st(\mathcal{F})$ are down sets and both families are monotone, meaning that if $\mathcal{F} \subseteq \mathcal{F}'$ are set systems then $Sh(\mathcal{F}) \subseteq Sh(\mathcal{F}')$ and $st(\mathcal{F}) \subseteq st(\mathcal{F}')$. For the size of $st(\mathcal{F})$ one can prove the so called reverse Sauer inequality:

Proposition 2.1 *(see [3]) $|st(\mathcal{F})| \leq |\mathcal{F}|$ for every set system $\mathcal{F} \subseteq 2^{[n]}$.*

Bollobás and Radcliffe in [4] obtained several important results concerning shattering and strong shattering, including:

Proposition 2.2 *(see [4], Theorem 2) $\mathcal{F} \subseteq 2^{[n]}$ is extremal with respect to the Sauer inequality (i.e. is shattering-extremal) iff it is extremal with respect to the reverse Sauer inequality i.e. $|st(\mathcal{F})| = |\mathcal{F}| \iff |Sh(\mathcal{F})| = |\mathcal{F}|$.*

Since the two extremal cases coincide, we will call such set systems shortly just extremal. As a consequence of the above facts, we get, that for extremal systems we have $st(\mathcal{F}) = Sh(\mathcal{F})$.

For $i \in [n]$ let φ_i be the ith *bit flip operation*, i.e. for $F \in 2^{[n]}$ we have

$$\varphi_i(F) = F \triangle \{i\} = \begin{cases} F \backslash \{i\} & \text{if } i \in F \\ F \cup \{i\} & \text{if } i \notin F \end{cases}$$

and $\varphi_i(\mathcal{F}) = \{\varphi_i(F) \mid F \in \mathcal{F}\}$. The family of shattered sets is trivially invariant under the bit flip operation, i.e. $Sh(\mathcal{F}) = Sh(\varphi_i(\mathcal{F}))$ for all $i \in [n]$, and hence so is extremality. This means that when dealing with a nonempty set system $\mathcal{F}$, and examining its extremality, we can assume that $\emptyset \in \mathcal{F}$, otherwise we could apply bit flips to it, to bring $\emptyset$ inside.

In terms of the inclusion graph, φ_i flips the directions of edges with label i, i.e. there is a bijection between the vertices of $\mathbb{G}_\mathcal{F}$ and $\mathbb{G}_{\varphi_i(\mathcal{F})}$ that preserves all edges with label different from i, and reverses edges with label i. This bijection is simply given by the reflection with respect to the hyperplane $x_i = \frac{1}{2}$ in the Hamming graph, when viewed as a subset of $\mathbb{R}^n$.

Note that for any set system $\mathcal{F} \subseteq 2^{[n]}$, the identity map naturally embeds the inclusion graph $\mathbb{G}_\mathcal{F}$ into $\mathbb{H}_n$. We say that the inclusion graph $\mathbb{G}_\mathcal{F}$ is

isometrically embedded (into $\mathbb{H}_n$), if this embedding is an isometry, meaning that for arbitrary $F, G \in \mathcal{F}$ we have $d_{\mathbb{G}_\mathcal{F}}(F, G) = d_{\mathbb{H}_n}(F, G)$, i.e. there is a path of length $d_{\mathbb{H}_n}(F, G) = |F \triangle G|$ between F and G inside the undirected version of $\mathbb{G}_\mathcal{F}$. Greco in [7] proved the following:

Proposition 2.3 *If $\mathcal{F} \subseteq 2^{[n]}$ is extremal, then $\mathbb{G}_\mathcal{F}$ is isometrically embedded.*

As this fact will be used several times, we provide the reader with a simple proof from [9]:

Proof: Suppose the contrary, namely that $\mathbb{G}_\mathcal{F}$ is not isometrically embedded. Then there exist sets $A, B \in \mathcal{F}$ such that $d_{\mathbb{H}_n}(A, B) = k < d_{\mathbb{G}_\mathcal{F}}(A, B)$. Suppose that A and B are such that k is minimal. Clearly $k \geq 2$. W.l.o.g we may suppose that $A = \emptyset$ and $|B| = k$, otherwise one could apply bit flips to the set system to achieve this. Note that distances both in $\mathbb{G}_\mathcal{F}$ and in $\mathbb{H}_n$ are invariant under bit flips.

We claim that there is no set $C \in \mathcal{F}$ different from A with $C \subsetneq B$. Indeed suppose such C exists, then

$$d_{\mathbb{H}_n}(A, C) + d_{\mathbb{H}_n}(C, B) = d_{\mathbb{H}_n}(A, B) = k < d_{\mathbb{G}_\mathcal{F}}(A, B) \leq d_{\mathbb{G}_\mathcal{F}}(A, C) + d_{\mathbb{G}_\mathcal{F}}(C, B).$$

From this we have either $d_{\mathbb{H}_n}(A, C) < d_{\mathbb{G}_\mathcal{F}}(A, C)$ or $d_{\mathbb{H}_n}(C, B) < d_{\mathbb{G}_\mathcal{F}}(C, B)$. Since $d_{\mathbb{H}_n}(A, C), d_{\mathbb{H}_n}(C, B) < k$ we get a contradiction in both cases with the minimality of k.

Now since $\mathcal{F}$ is extremal, so must be $\mathcal{F}_{\emptyset, B}$. However in our case $\mathcal{F}_{\emptyset, B} = \{\emptyset, B\}$, and so if $B = \{b_1, \ldots, b_k\}$, then $Sh(\mathcal{F}_{\emptyset, B}) = \{\emptyset, \{b_1\}, \ldots, \{b_k\}\}$. Counting cardinalities we get that $|Sh(\mathcal{F}_{\emptyset, B})| = |B| + 1 = k + 1 \geq 3 > 2 = |\mathcal{F}_{\emptyset, B}|$, implying that $\mathcal{F}_{\emptyset, B}$ cannot be extremal. This contradiction finishes the proof. ∎

It is easy to see that $S \in st(\mathcal{F})$ (and so in the extremal case $S \in Sh(\mathcal{F})$) is just equivalent to the fact that $\mathbb{G}_{2^S}$ is isomorphic to a subgraph of $\mathbb{G}_\mathcal{F}$ as a directed edge labelled graph, i.e. there exists a bijection between the vertices of $\mathbb{G}_{2^S}$ and $2^{|S|}$ vertices of $\mathbb{G}_\mathcal{F}$ preserving edges, edge labels and edge directions. If this happens, we will say, that there is a copy of $\mathbb{G}_{2^S}$ in $\mathbb{G}_\mathcal{F}$.

Suppose that for a set $S \subseteq [n]$ there are 2 different copies of $\mathbb{G}_{2^S}$ in $\mathbb{G}_\mathcal{F}$, i.e. there are two different sets $I_1, I_2 \subseteq [n] \backslash S$ such that $2^S + I_1, 2^S + I_2 \subseteq \mathcal{F}$. Since $I_1 \neq I_2$, there must be an element $\alpha \notin S$ such that $\alpha \in I_1 \triangle I_2$. For this element α we clearly have that $\mathcal{F}$ shatters $S \cup \{\alpha\}$.

Observation 2.1 *If $\mathcal{F} \subseteq 2^{[n]}$ is extremal and $S \subseteq [n]$ is a maximal element in $st(\mathcal{F}) = Sh(\mathcal{F})$, in the sense that $S \in st(\mathcal{F}) = Sh(\mathcal{F})$ and for all $S' \supsetneq S$ we have $S' \notin st(\mathcal{F}) = Sh(\mathcal{F})$, then S is uniquely strongly shattered, i.e. there is one unique copy of $\mathbb{G}_{2^S}$ in $\mathbb{G}_{\mathcal{F}}$*

Indeed, by the earlier reasoning, multiple copies would result a contradiction with the maximality of S.

3 Construction of extremal families

In this section we will describe and study a process for building up an extremal set system on the ground set $[n]$ together with its inclusion graph. First we describe the building process for the set system and then study how the inclusion graph evolves in the meantime. Let Step 0 be the initialization, after which we are given the set system $\{\emptyset\}$. Now suppose we are given a set system $\mathcal{F}$ and consider the following two types of operations to enlarge $\mathcal{F}$:

- **Step A** - If such exists, take an element $\alpha \in [n]\backslash supp(\mathcal{F})$ together with a set $W \in \mathcal{F}$ and add the set $V = \{W, \alpha\}$ to $\mathcal{F}$.

  Note that the singleton $\{\alpha\}$ is strongly shattered by $\mathcal{F} \cup \{V\}$, as shown by the sets W and V, but is not by $\mathcal{F}$, as by assumption $\alpha \notin supp(\mathcal{F})$.

- **Step B** - If there exist, take two elements $\alpha, \beta \in supp(\mathcal{F})$ such that $\{\alpha, \beta\} \notin st(\mathcal{F})$, together with sets $P, W, Q \in \mathcal{F}$ such that $Q \triangle W = \{\alpha\}$ and $P \triangle W = \{\beta\}$. Let $V = W \triangle \{\alpha, \beta\}$. V is also the unique set satisfying $P \triangle V = \{\alpha\}$ and $Q \triangle V = \{\beta\}$. For these sets we have that $\{P, W, Q, V\} = W \cap V + 2^{\{\alpha,\beta\}} = P \cap Q + 2^{\{\alpha,\beta\}}$, and hence V cannot belong to $\mathcal{F}$, otherwise the sets P, W, Q, V would strongly shatter $\{\alpha, \beta\}$, contradicting our assumption. Therefore, it is reasonable to add V to $\mathcal{F}$.

  Note that the set $\{\alpha, \beta\}$ is strongly shattered by $\mathcal{F} \cup \{V\}$, as shown by the sets P, W, Q and V, but is not by $\mathcal{F}$ by assumption.

Let $\mathcal{E}$ be the collection of all set systems $\mathcal{F}$ that can be built up starting with Step 0 and then using steps of type A and B in an arbitrary but valid order.

Lemma 3.1 *Any set system $\mathcal{F} \in \mathcal{E}$ is extremal and $dim_{VC}(\mathcal{F}) \leq 2$.*

Proof: We will use induction on the size of $\mathcal{F}$. If $|\mathcal{F}| = 1$ then necessarily $\mathcal{F} = \{\emptyset\}$, which is clearly extremal and $dim_{VC}(\mathcal{F}) = 0$. Now suppose we know the result for all members of $\mathcal{E}$ of size at most $m \geq 1$, and consider a system $\mathcal{F} \in \mathcal{E}$ of size $m + 1$. As $\mathcal{F} \in \mathcal{E}$ it can be built up starting from $\{\emptyset\}$ using Steps A and B. Fix one such building process, and let $\mathcal{F}'$ be the set system before the last building step. As noted previously, independently of the type of the last step there is a set S that is strongly shattered by $\mathcal{F}$ but is not strongly shattered by $\mathcal{F}'$. S is either a singleton or a set of size 2, depending on the type of the last step. By the induction hypothesis $\mathcal{F}'$ is extremal and $dim_{VC}(\mathcal{F}') \leq 2$. Using the reverse Sauer inequality we get that

$$|\mathcal{F}'| = |st(\mathcal{F}')| < |st(\mathcal{F})| \leq |\mathcal{F}| = |\mathcal{F}'| + 1,$$

what is possible only if $|st(\mathcal{F})| = |st(\mathcal{F}')| + 1$ and $|st(\mathcal{F})| = |\mathcal{F}|$, in particular $\mathcal{F}$ is extremal.

However in the extremal case the family of shattered sets is the same as the family of strongly shattered sets, and so the above reasoning also gives that there is exactly one set that is shattered by $\mathcal{F}$ and is not shattered by $\mathcal{F}'$, namely S, and so $dim_{VC}(\mathcal{F}) \leq max(dim_{VC}(\mathcal{F}'), |S|) \leq 2$. ∎

The proof of Lemma 3.1 also describes how the family of shattered/strongly shattered sets grows during a building process. After each step it grows by exactly one new set, namely by $\{\alpha\}$, if the step considered was Step A with the label α, and by $\{\alpha, \beta\}$, if the step considered was Step B with labels α, β. By our assumptions on the steps it also follows that a valid building process for a set system $\mathcal{F} \in \mathcal{E}$ cannot involve twice Step A with the same label α, neither twice Step B with the same pair of labels α, β, and we also have that

$$Sh(\mathcal{F}) = st(\mathcal{F}) = \begin{cases} \{\emptyset\} \cup \{\{\alpha\} \mid \text{Step A is used with label } \alpha\} \cup \\ \{\{\alpha, \beta\} \mid \text{Step B is used with labels } \alpha \text{ and } \beta\} \end{cases}.$$

Now consider a valid building process from $\mathcal{E}$, and let us examine, how the inclusion graph evolves. We use the notation from the definitions of Steps A and B. Suppose we have already built up a set system $\mathcal{F}$, and we are given its inclusion graph $G_\mathcal{F}$.

In Step A we add a new vertex, namely V to $G_\mathcal{F}$, together with one new directed edge with label α going from W to V. As $\alpha \notin supp(\mathcal{F})$, V has no other neighbors in $G_\mathcal{F}$. Figure 1 shows Step A in terms of the inclusion graph.
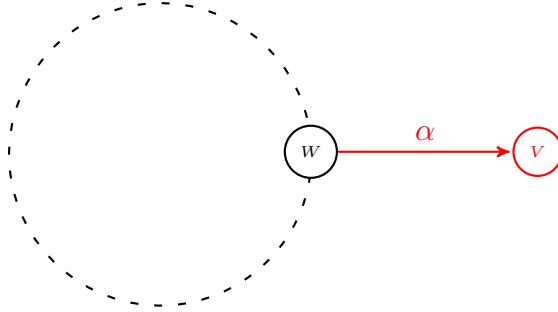
Figure 1: Step A

In Step B we also add one new vertex to $G_{\mathcal{F}}$, namely V. As the distance of V from both P and Q is 1, and $P \triangle V = \{\alpha\}$ and $Q \triangle V = \{\beta\}$, we have to add at least 2 new edges, one between P and V with label α and one between Q and V with label β. The direction of these edges is predetermined by the vertices P, W and Q. Figure 2 shows all possible cases for the directions of these edges. We claim that no other edges need to be added, i.e. V has no other neighbors in $G_{\mathcal{F}}$. Indeed suppose that the new vertex V has another neighbor X in $G_{\mathcal{F}}$, different from P and Q, that should be connected to it with some label γ different from α and β. See Figure 3, where edge directions are ignored, only edge labels are shown.

Here $d_{\mathbb{H}_n}(P, X) = |P \triangle X| = |\{\alpha, \gamma\}| = 2$. On the other hand as $\mathcal{F}$ was built using Steps A and B starting from $\{\emptyset\}$, it is a member of $\mathcal{E}$, and so by Lemma 8 it is extremal. According to Proposition 2.3 this implies that $\mathcal{G}_{\mathcal{F}}$ is isometrically embedded. This means that there should be a vertex Y in $\mathbb{G}_{\mathcal{F}}$ connected to both P and X with edges with labels γ and α respectively. The same reasoning applies for Q and V with some intermediate vertex Z and edge labels β, γ. However in this case, independently of the directions of the edges, we have $\{X \cap \{\alpha, \beta\}, Y \cap \{\alpha, \beta\}, Z \cap \{\alpha, \beta\}, W \cap \{\alpha, \beta\}\} = 2^{\{\alpha, \beta\}}$, i.e. the sets X, Y, Z, W shatter the set $\{\alpha, \beta\}$, and so by the extremality of $\mathcal{F}$ we have that $\{\alpha, \beta\}$ is also strongly shattered, what contradicts the assumptions of Step B.

From now on it will depend on the context whether we regard Steps A and B as building steps for extremal set systems of VC dimension at most 2 or as building steps for their inclusion graphs.
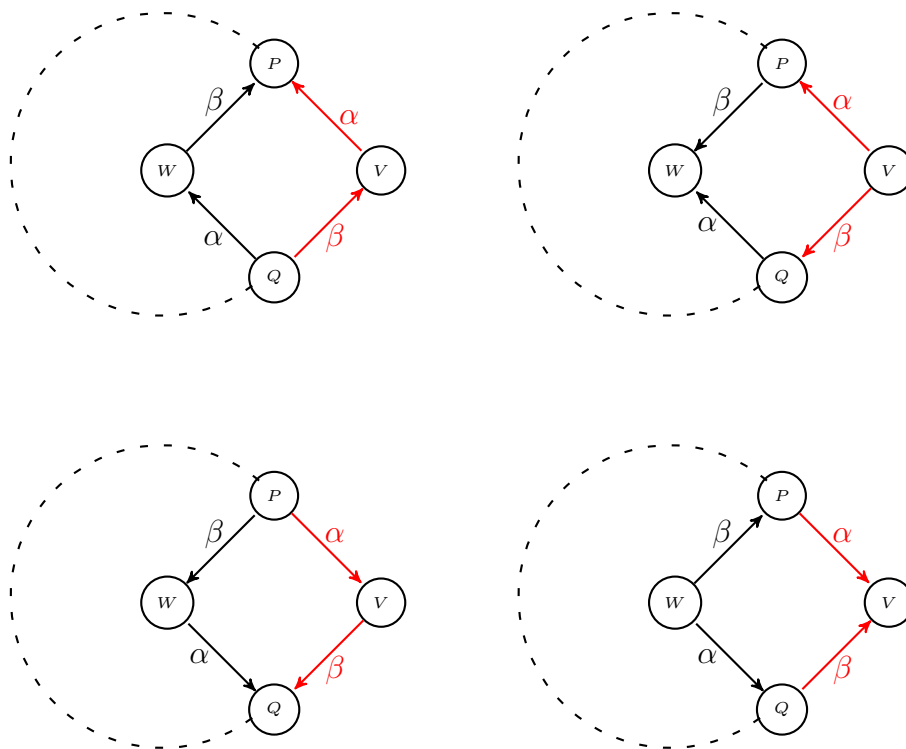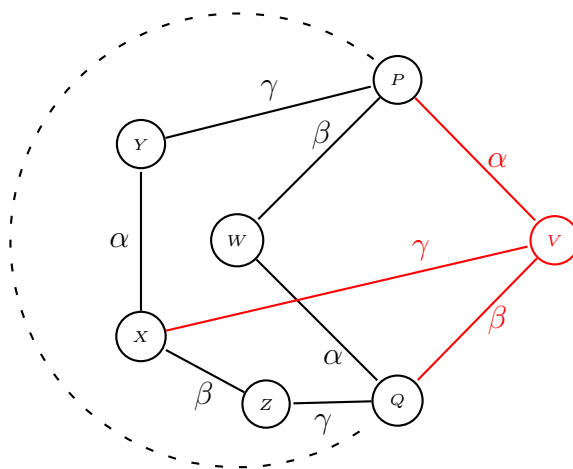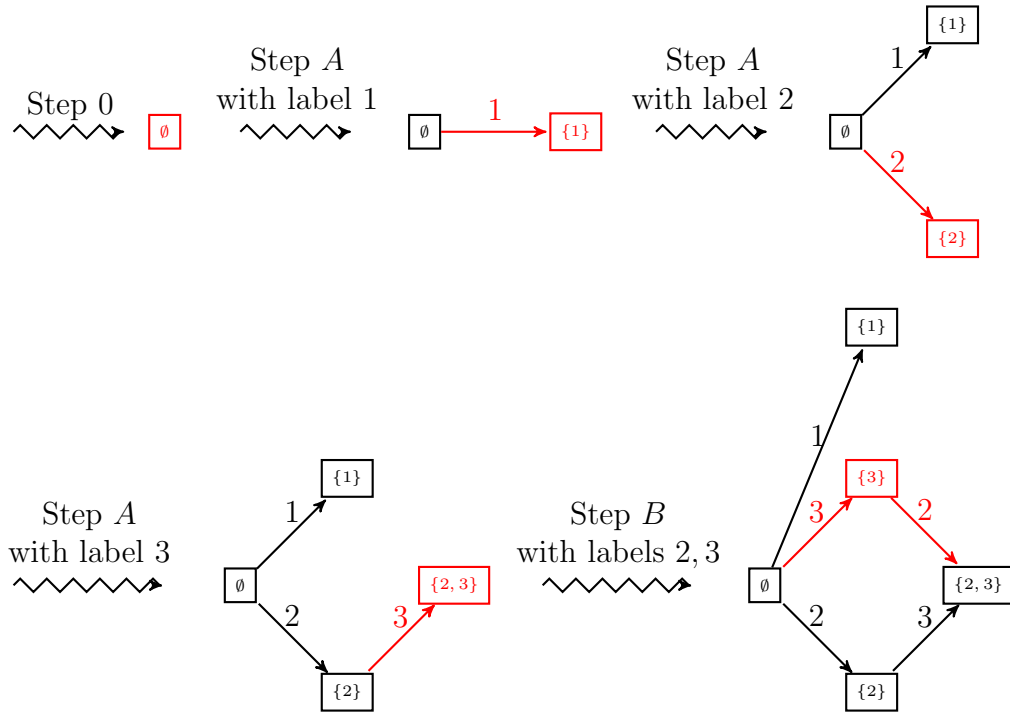
Figure 2: Step B

Figure 3: Case of Step B

Figure 4: Example of a building process in $\mathcal{E}$

Figure 4 shows a possible building process in $\mathcal{E}$ for the set system

$$\mathcal{F} = \{\emptyset, \{1\}, \{2\}, \{3\}, \{2, 3\}\}$$

in terms of the inclusion graph.

Take an element of $\mathcal{E}$ and fix a valid building process for it. The above observations also imply, that when observing the evolution of the inclusion graph, after the first occurrence of an edge with some fixed label α, new edges with the same label can come up only when using Step B always with a different label next to α. By easy induction on the number of building steps, this results that between any two edges with the same label α there is a "path of 4-cycles". See Figure 5. Note that as in Figure 5, all the β_i's must be different. Along this path of 4-cycles we also obtain a shortest path between X_1 and X_2, and similarly between Y_1 and Y_2.
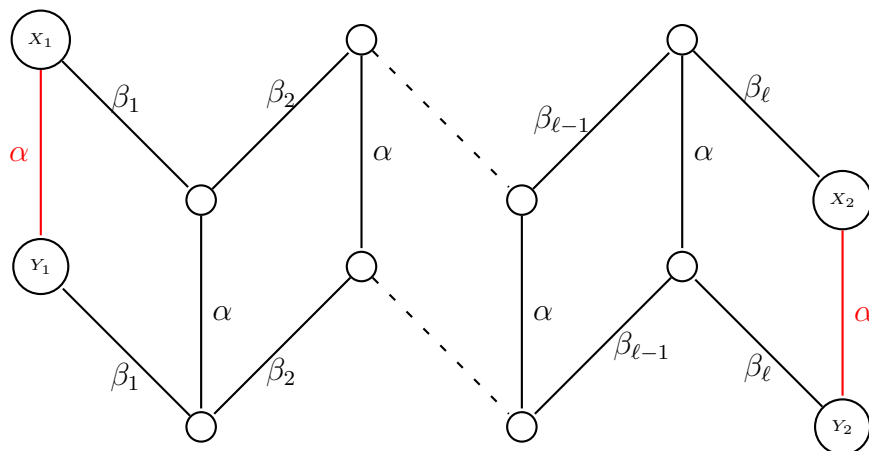
Figure 5: Path of 4 cycles

4 Main results

The first of the main results of this paper is that the set systems in $\mathcal{E}$, described in the previous section, are actually all the extremal set systems of VC-dimension at most 2 and containing $\emptyset$.

Theorem 4.1 $\mathcal{F} \subseteq 2^{[n]}$ *is an extremal set system with* $dim_{VC}(\mathcal{F}) \leq 2$ *and* $\emptyset \in \mathcal{F}$ *iff* $\mathcal{F} \in \mathcal{E}$.

Before turning to the proof of Theorem 4.1, we first prove a lemma about the building processes in $\mathcal{E}$, that will play a key role further on.

Lemma 4.1 *Suppose that* $\mathcal{F}', \mathcal{F}$ *are elements of* $\mathcal{E}$ *such that* $\mathcal{F}' \subseteq \mathcal{F}$. *Then* $\mathcal{F}'$ *can be extended with valid building process to build up* $\mathcal{F}$.

Proof: Suppose this is not the case, and consider a counterexample. Without loss of generality we may suppose that the counterexample is such that $\mathcal{F}'$ cannot be continued with any valid step towards $\mathcal{F}$. $\mathcal{F}'$ and $\mathcal{F}$ are both extremal and so $\mathbb{G}_{\mathcal{F}'}$ and $\mathbb{G}_{\mathcal{F}}$ are both isometrically embedded, in particular connected, hence the neighborhood of $\mathbb{G}_{\mathcal{F}'}$ inside $\mathbb{G}_{\mathcal{F}}$ is nonempty. Now take a closer look at the edges on the boundary of $\mathbb{G}_{\mathcal{F}'}$.

If there would be an edge going out from $\mathbb{G}_{\mathcal{F}'}$ with a label $\alpha \in supp(\mathcal{F}) \backslash supp(\mathcal{F}')$, then Step A would apply with this label α. On the other hand there cannot

be an edge going into $\mathbb{G}_{\mathcal{F}'}$ with a label $\alpha \notin supp(\mathcal{F}')$, otherwise the endpoint of this edge inside $\mathbb{G}_{\mathcal{F}'}$ would contain α, what would be a contradiction.

We can therefore assume that the label of any edge on the boundary of $\mathbb{G}_{\mathcal{F}'}$, independently of the direction of the edge, is an element of $supp(\mathcal{F}')$. However as $\emptyset \in \mathcal{F}'$ and $\mathbb{G}_{\mathcal{F}'}$ is isometrically embedded, an element belongs to $supp(\mathcal{F}')$ only if it appears as an edge label in $\mathbb{G}_{\mathcal{F}'}$. Now take an edge (W, V) on the boundary of $\mathbb{G}_{\mathcal{F}'}$ with $W \in \mathcal{F}'$, $V \in \mathcal{F}\backslash\mathcal{F}'$ and with some label α, together with an edge (X, Y) with the same label inside $\mathbb{G}_{\mathcal{F}'}$. Denote the distance of the edges (W, V) and (X, Y) by ℓ, i.e. $d_{\mathbb{H}_n}(W, X) = d_{\mathbb{H}_n}(V, Y) = \ell$. The latter equality means, that depending on the direction of the edges, W and X both do contain the element α, or neither of them does. Suppose that the triple $\alpha, (W, V), (X, Y)$ is such that the distance ℓ is minimal.

First suppose that $\ell > 1$. Since the edges $(W, V), (X, Y)$ have the same label and $\mathcal{F} \in \mathcal{E}$, there is a path of 4-cycles of length ℓ between them inside $\mathbb{G}_{\mathcal{F}}$. This path of 4-cycles also provides shortest paths between the endpoints of the edges $(W, V), (X, Y)$. By the minimality of our choice, in this path, except the edges at the ends, there cannot be an edge with label α neither totally inside $\mathbb{G}_{\mathcal{F}'}$, neither on the boundary of it, meaning that this path of 4-cycles is essentially going outside $\mathbb{G}_{\mathcal{F}'}$. See Figure 6.

Since $\mathbb{G}_{\mathcal{F}'}$ is isometrically embedded and $d_{\mathbb{H}_n}(W, X) = \ell$, there must be a path of length ℓ between W and X inside $\mathbb{G}_{\mathcal{F}'}$. As this path runs inside $\mathbb{G}_{\mathcal{F}'}$, it has to be disjoint from the path of 4-cycles. Along the path of 4-cycles all the β_i's are different, so for each i exactly one of the sets W and X contains the element β_i. In particular for $i = 1$, the shortest path between W and X inside $\mathbb{G}_{\mathcal{F}'}$ also has to contain an edge (T, S) with label β_1 with direction determined by the sets W and X. However the distance between W and T is at most $\ell - 1$, and hence the triple $\beta_1, (W, Q_1), (T, S)$ contradicts with the minimality of the initial triple $\alpha, (W, V), (X, Y)$ where the distance was ℓ.

By the above reasoning only $\ell = 1$ is possible. In this case the endpoints of the edges (W, V), (X, Y) are connected by edges with the same label. Let this label be β. The direction of these edges is predetermined by $\mathbb{G}_{\mathcal{F}'}$. $\{\alpha, \beta\} \notin st(\mathcal{F}')$, otherwise there would be already a copy of $\mathbb{G}_{2^{\{\alpha,\beta\}}}$ in $\mathbb{G}_{\mathcal{F}'}$, which together with the vertices W, V, X, Y would give us two different copies of it inside $\mathbb{G}_{\mathcal{F}}$, which is impossible by Observation 2.1, as $\{\alpha, \beta\}$ is a maximal set strongly shattered by the extremal family $\mathcal{F}$. Hence Step B applies with new vertex V, edges $(W, V), (V, Y)$ and labels α, β respectively, contradicting with the fact, that we started with a counterexample. See Figure 7. $\blacksquare$
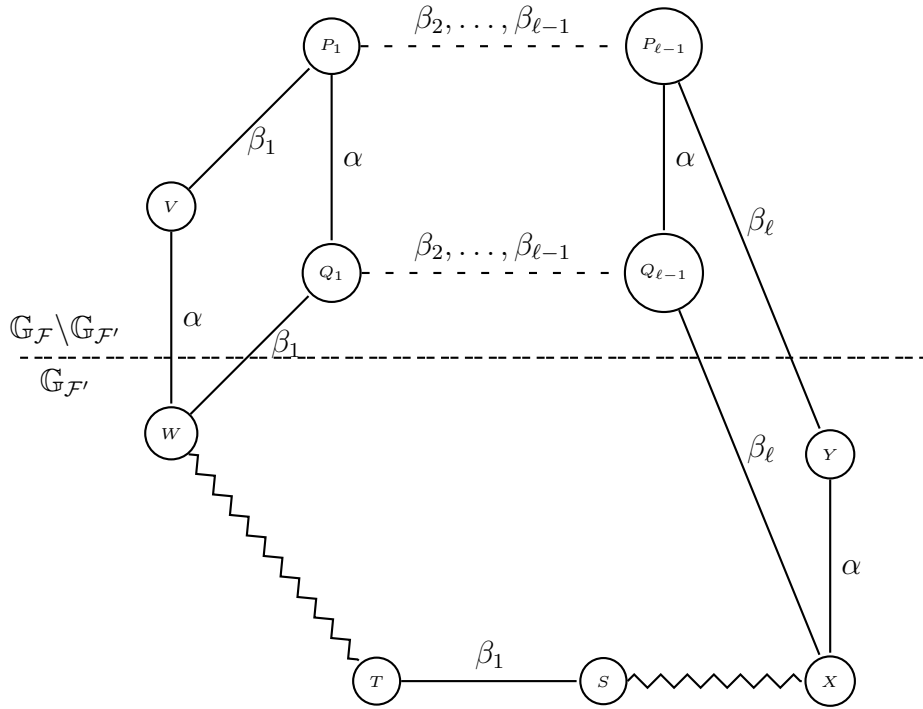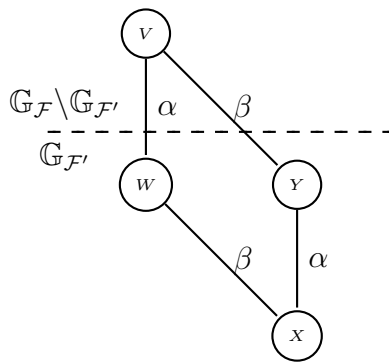
13

Figure 6: Case $l > 1$

Figure 7: Case $\ell = 1$

Now we are ready to prove Theorem 4.1.

Proof: One direction of the theorem is just Lemma 3.1. For the other direction we use induction on the number of sets in $\mathcal{F}$. If $|\mathcal{F}| = 1$, then $\mathcal{F}$ is necessarily $\{\emptyset\}$, and so belongs trivially to $\mathcal{E}$. Now suppose we proved the statement for all set systems with at most $m - 1$ members, and let $\mathcal{F}$ be an extremal family of size m, of VC-dimension at most 2 and containing $\emptyset$. Take an arbitrary element α appearing as a label of an edge going out from $\emptyset$ in $\mathbb{G}_{\mathcal{F}}$, i.e. an element α such that $\{\alpha\} \in \mathcal{F}$. Consider the standard subdivision of $\mathcal{F}$ with respect to the element α with parts $\mathcal{F}_0$ and $\mathcal{F}_1$ (see Definition 2.1), and let

$$\widehat{\mathcal{F}}_1 = \{F \cup \{\alpha\} \; : \; F \in \mathcal{F}_1\}.$$

Note that with respect to shattering and strong shattering $\mathcal{F}_1$ and $\widehat{\mathcal{F}}_1$ behave in the same way. Since $\mathcal{F}$ is extremal, so are $\mathcal{F}_0$, $\mathcal{F}_1$ and hence $\widehat{\mathcal{F}}_1$ as well, and clearly their VC-dimension is at most 2. The collection of all edges with label α in the inclusion graph $\mathbb{G}_{\mathcal{F}}$ forms a cut. This cut divides $\mathbb{G}_{\mathcal{F}}$ into two parts, that are actually the inclusion graphs $\mathbb{G}_{\mathcal{F}_0}$ and $\mathbb{G}_{\widehat{\mathcal{F}}_1}$. Note that $\mathbb{G}_{\mathcal{F}_1}$ and $\mathbb{G}_{\widehat{\mathcal{F}}_1}$ are isomorphic as directed edge labelled graphs. Let T_0 and T_1 be the induced subgraphs on the endpoints of the cut edges in $\mathbb{G}_{\mathcal{F}_0}$ and $\mathbb{G}_{\widehat{\mathcal{F}}_1}$, respectively. See Figure 8. T_0 and T_1 are isomorphic, and they are actually the inclusion graphs of the set systems $\mathcal{T}_0 = \mathcal{F}_0 \cap \mathcal{F}_1$ and $\mathcal{T}_1 = \{F \cup \{\alpha\}, F \in \mathcal{T}_0\}$. Similarly to the pair $\mathcal{F}_1, \widehat{\mathcal{F}}_1$, the set systems $\mathcal{T}_0$ and $\mathcal{T}_1$ also behave in the same way with respect to shattering and strong shattering. By assumption $\mathcal{F}$ is extremal, and so according to Proposition 5.1 from [11] so is $\mathcal{T}_0$ and hence $\mathcal{T}_1$. For every set S in $Sh(\mathcal{T}_0) = Sh(\mathcal{F}_0 \cap \mathcal{F}_1) \subseteq 2^{[n] \setminus \{\alpha\}}$ the set $S \cup \{\alpha\}$ is shattered by $\mathcal{F}$, implying that $dim_{VC}(\mathcal{T}_0) \leq dim_{VC}(\mathcal{F}) - 1 \leq 1$. Therefore $\mathcal{T}_0$ is an extremal family of VC-dimension at most 1, and so by Proposition 1.2 we get that T_0 (and hence T_1) is a directed edge labelled tree having all edge labels different. Note that for any edge label β appearing in T_0 (and hence in T_1), there is a copy of $\mathbb{G}_{2^{\{\alpha,\beta\}}}$ along the cut, implying that $\{\alpha, \beta\} \in st(\mathcal{F}) = Sh(\mathcal{F})$. By the VC-dimension constraint on $\mathcal{F}$ the set $\{\alpha, \beta\}$ is a maximal element of $st(\mathcal{F}) = Sh(\mathcal{F})$, and so by Observation 2.1 there cannot be another copy of $\mathbb{G}_{2^{\{\alpha,\beta\}}}$ in $\mathbb{G}_{\mathcal{F}}$, neither in $\mathbb{G}_{\mathcal{F}_0}$ nor in $\mathbb{G}_{\widehat{\mathcal{F}}_1}$, in particular $\{\alpha, \beta\} \notin st(\mathcal{F}_0)$.

Let's now turn to the building process of $\mathcal{F}$. Our choice of α guarantees that $\emptyset \in \mathcal{F}_0, \mathcal{F}_1$ and so by the induction hypothesis both of them belong to $\mathcal{E}$. In particular we can build up $\mathcal{F}_0$, and in the meantime $\mathbb{G}_{\mathcal{F}_0}$, according to the
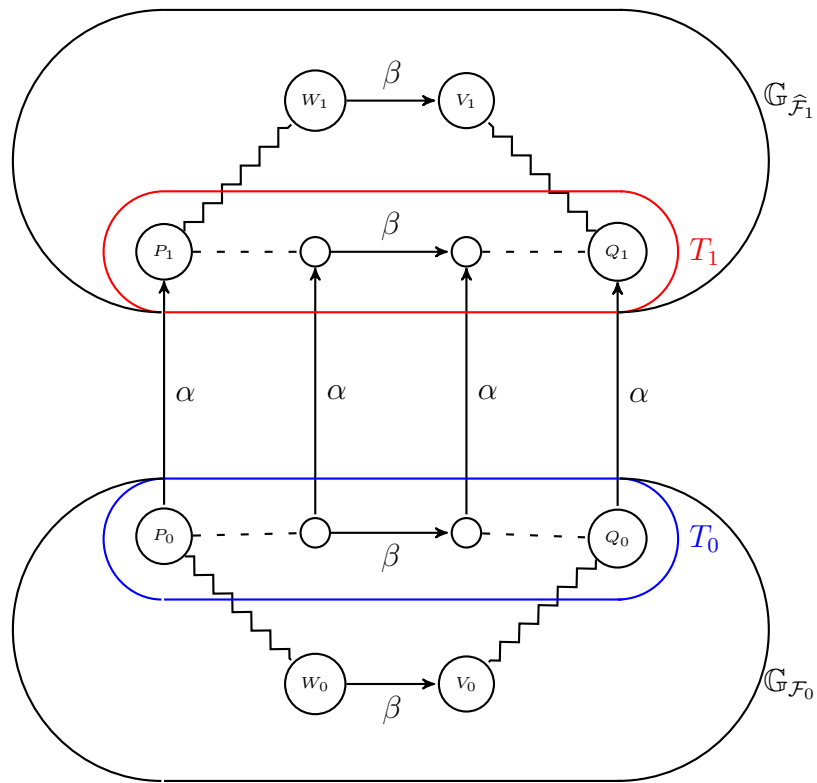
Figure 8: Building up extremal set systems

16

building rules in $\mathcal{E}$. $\alpha \notin supp(\mathcal{F}_0)$ and so we can apply Step A with α to add one fixed cut edge to $\mathbb{G}_{\mathcal{F}_0}$. Then we apply Step B several times to add the whole of T_1 to $\mathbb{G}_{\mathcal{F}_0}$ and simultaneously $\mathcal{T}_1$ to $\mathcal{F}_0$. By earlier observations all edge labels of T_1 are different, and if β is such a label, then $\{\alpha, \beta\} \notin st(\mathcal{F}_0)$, and hence all these applications of Step B will be valid ones. The building process so far shows that $\mathcal{F}_0 \cup \mathcal{T}_1$ is also a member of $\mathcal{E}$. $\mathbb{G}_{\mathcal{F}_0 \cup \mathcal{T}_1}$ is just $\mathbb{G}_{\mathcal{F}_0}$ and T_1 glued together along the cut in the way described above.

T_0 shows that $\mathcal{T}_0$ can be built up using only Step A, and hence it belongs to $\mathcal{E}$. The inclusion $\mathcal{T}_1 \subseteq \widehat{\mathcal{F}}_1$ shows that $\mathcal{T}_0 \subseteq \mathcal{F}_1$, therefore by Lemma 4.1 $\mathcal{T}_0$ can be extended with a valid building process to build up $\mathcal{F}_1$. This extension can also be considered as building up $\widehat{\mathcal{F}}_1$ from $\mathcal{T}_1$. $\emptyset \notin \mathcal{T}_1, \widehat{\mathcal{F}}_1$ and so neither of the two systems is a member of $\mathcal{E}$, however this causes no problems, as the pairs $\mathcal{T}_0$, $\mathcal{T}_1$ and $\mathcal{F}_1$, $\widehat{\mathcal{F}}_1$ behave in the same way with respect to shattering and strong shattering, and so all building steps remain valid.

We claim, that this last building procedure remains valid, and so completes a desired building process for $\mathcal{F}$, if we start from $\mathcal{F}_0 \cup \mathcal{T}_1$ instead of $\mathcal{T}_1$. First note that if there is a label appearing both in $\mathbb{G}_{\mathcal{F}_0}$ and $\mathbb{G}_{\widehat{\mathcal{F}}_1}$, then it appears also in T_0, and hence in T_1. Indeed let β be such a label, and consider 2 edges with this label, one going from W_0 to V_0 in $\mathbb{G}_{\mathcal{F}_0}$ and the other going from W_1 to V_1 in $\mathbb{G}_{\widehat{\mathcal{F}}_1}$. See Figure 8. $\mathbb{G}_{\mathcal{F}}$ is isometrically embedded, therefor there is a shortest path both between W_0 and W_1 and between V_0 and V_1 in $\mathbb{G}_{\mathcal{F}}$. Thanks to β these two paths have to be disjoint. Both of these paths must have a common edge with the cut, say (P_0, P_1) and (Q_0, Q_1), with P_0 and Q_0 in $\mathbb{G}_{\mathcal{F}_0}$. Since $\beta \in P_0 \triangle Q_0$, along the shortest path between P_0 and Q_0 in the isometrically embedded inclusion graph T_0 of the extremal family $\mathcal{T}_0$ there must be an edge with label β. According to this, when applying Step A in the extension process, then the used element will be new not just when we start from $\mathcal{T}_1$, but also when starting from $\mathcal{F}_0 \cup \mathcal{T}_1$.

Finally suppose that an application of Step B with some labels β, γ in the extension process turns invalid when we start from $\mathcal{F}_0 \cup \mathcal{T}_1$ instead of $\mathcal{T}_1$. This is possible only if $\{\beta, \gamma\} \in st(\mathcal{F}_0 \cup \mathcal{T}_1) \backslash st(\mathcal{T}_0)$, i.e. there is a copy of $\mathbb{G}_{2\{\beta, \gamma\}}$ already in $\mathbb{G}_{\mathcal{F}_0 \cup \mathcal{T}_1}$. However this copy together with the copy, that the invalid use of Step B results, gives two different occurrences of $\mathbb{G}_{2\{\beta, \gamma\}}$ inside $\mathbb{G}_{\mathcal{F}}$, which is impossible by Observation 2.1, as $\{\beta, \gamma\}$ is a maximal set strongly shattered by the extremal family $\mathcal{F}$. ∎

As a corollary of Theorem 4.1 one can solve an open problem, posed in [11], in the special case when the VC-dimension of the systems investigated

is bounded by 2.

Open problem 1 *(See [11]) For a nonempty s-extremal family $\mathcal{F} \subseteq 2^{[n]}$ does there always exist a set $F \in \mathcal{F}$ such that $\mathcal{F} \backslash \{F\}$ is still s-extremal?*

The case when the VC-dimension of the systems investigated is bounded by 1 was solved in [11]. Here we propose a solution for set systems of VC-dimension at most 2.

Theorem 4.2 *Let $\mathcal{F} \subseteq 2^{[n]}$ be a nonempty extremal family of VC dimension at most 2. Then there exists an element $F \in \mathcal{F}$ such that $\mathcal{F} \backslash \{F\}$ is still extremal.*

Proof: Let $F \in \mathcal{F}$ be an arbitrary set from the set system. Recall that φ_i is the ith bit flip operation, and let $\varphi = \prod_{i \in F} \varphi_i$. Since bit flips preserve extremality, $\varphi(\mathcal{F})$ is extremal as well. Moreover $\varphi(F) = \emptyset \in \varphi(\mathcal{F})$, and so by Theorem 4.1 we have $\varphi(\mathcal{F}) \in \mathcal{E}$, hence we can consider a building process for it. Let $V \in \varphi(\mathcal{F})$ be the set added in the last step of this building process. The same building process shows that $\varphi(\mathcal{F}) \backslash \{V\} \in \mathcal{E}$, and hence by Theorem 4.1 we have that $\varphi(\mathcal{F}) \backslash \{V\}$ is an extremal family of VC dimension at most 2 and containing $\emptyset$. However $\varphi(\mathcal{F}) \backslash \{V\} = \varphi(\mathcal{F} \backslash \{\varphi(V)\})$, and since bit flips preserve extremality, we get that $\varphi(\varphi(\mathcal{F} \backslash \{\varphi(V)\})) = \mathcal{F} \backslash \{\varphi(V)\}$ is also extremal, meaning that the set $\varphi(V) \in \mathcal{F}$ can be removed from the extremal system $\mathcal{F}$ so that the result is still extremal. ■

5 Concluding remarks and future work

The building process from Section 2 can be generalized to the case when the VC-dimension bound is some fixed natural number t. We can define a building step for every set $S \subseteq [n]$ with $|S| \leq t$. Let Step($\emptyset$) be the initialization, after which we are given the set system $\{\emptyset\}$. For some set $S \subseteq [n]$ with $|S| \leq t$, Step(S) can be applied to a set system $\mathcal{F}$, if there exists some set $F \subseteq [n]$, $F \notin \mathcal{F}$, such that $S \in st(\mathcal{F} \cup \{F\}) \backslash st(\mathcal{F})$. If such set F exists, choose one, and let the resulting system be $\mathcal{F} \cup \{F\}$. In terms of the inclusion graph $S \in st(\mathcal{F} \cup \{F\}) \backslash st(\mathcal{F})$ means, that by adding the set F there arises a copy of $\mathbb{G}_{2^S}$ inside $\mathbb{G}_{\mathcal{F} \cup \{\mathcal{F}\}}$ containing the vertex F. Similarly as previously, one can prove that F's only neighbors are the ones contained

in this copy of $\mathbb{G}_{2^s}$. Using this observation $\text{Step}(S)$ could have been defined in terms of the inclusion graph as well (as it was done in the case $t = 2$).

Restrict our attention to those set systems, that can be built up starting with $\text{Step}(\emptyset)$, and then using always new building steps, i.e. not using a building step with the same set S twice. Along the same lines of thinking as in the case $t = 2$, one can prove that every such set system is extremal. We think, that these set systems are actually all the extremal families of VC-dimension at most t. Unfortunately, for the time being we were unable to prove a suitable generalization of Lemma 4.1. Once it is done, the generalization of Theorem 4.1, and as a corollary a generalization of Theorem 4.2 would follow easily. Although the general version Theorem 4.1 would not give such a transparent structural description of extremal systems as in the case $t = 1$, but still, its corollary, the generalization of Theorem 4.2 would solve the open problem proposed in [11] in its entire generality.

Acknowledgements

References

[1] R. Aharoni, R. Holzman, Personal communication, cited in [7]

[2] R.P. Anstee, L. Rónyai, A. Sali, Shattering News, Graphs and Combinatorics, Vol.18, 59-73 (2002)

[3] B. Bollobás, I. Leader, A.J. Radcliffe, Reverse Kleitman Inequalities, Proceedings of the London Mathematical Society, Vol. s3-58, 153-168 (1989)

[4] B. Bollobás, A.J. Radcliffe, Defect Sauer Results, Journal of Combinatorial Theory Series A, Vol. 72, 189-208 (1995)

[5] P. Frankl, S-extremal set systems, Handbook of combinatorics (Vol. 2), MIT Press, Cambridge, MA, 1996

[6] Z. Füredi, F. Quinn, Traces of Finite Sets, Ars Combinatoria, Vol. 18, 195-200 (1983)

[7] G. Greco, Embeddings and trace of finite sets, Information Processing Letters, Vol. 67, 199-203 (1998)

[8] L. Kozma, S. Moran, Shattering, Graph Orientations and Connectivity, Electronic Journal of Combintaorics, Vol. 20 - 3 (2013)

[9] T. Mészáros, S-extremal set systems and Gröbner bases, Diploma thesis, Budapest University of Technology and Economics (2010)

[10] L. Rónyai, T. Mészáros: Some Combinatorial Applications of Gröbner bases, Proc.CAI 2011, Lecture Notes in Computer Science, Vol. 6742

[11] T. Mészáros, L. Rónyai, Shattering extremal set systems of small VC-dimension, ISRN Combinatoric, Vol. 2013 (2013)

[12] A. Pajor, Sous-spaces 1: des Espaces de Banach, Travaux en Cours, Hermann, Paris, (1985)

[13] N. Sauer, On the Density of Families of Sets, Journal of Combinatorial Theory, Series A, Vol. 13, 145-147 (1972)

[14] S. Shelah, A Combinatorial Problem: Stability and Order for Models and Theories in Infinitary Language, Pacific Journal of Mathematics, Vol. 41, 247-261 (1972)