MATHEMATICS OF OPERATIONS RESEARCH
Vol. 30, No. 4, November 2005, pp. 880–896
ISSN 0364-765X | EISSN 1526-5471 | 05 | 3004 | 0880
DOI 10.1287/moor.1050.0161

The Number of Solutions Sufficient for Solving a Family of Problems

Ori Einstein
Department of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel,
einstei@post.tau.ac.il

Refael Hassin
Department of Statistics and Operations Research, Tel Aviv University, Tel Aviv 69978, Israel,
hassin@post.tau.ac.il, http://www.math.tau.ac.il/~hassin

This paper deals with families of optimization problems defined over a common set of potential solutions. We consider several problems-solutions systems, and for each one, prove the existence of a small set of solutions that contains an optimal solution to every problem. These proofs are mostly algebraic in nature. The families of problems covered here mostly include separation problems, problems on graphs and hypergraphs, and SAT problems.

Key words: hypergraphs; algebraic combinatorics; family of problems; cut tree; extremal combinatorics
MSC2000 subject classification: Primary: 05C15, 05C35, 05C40, 05C50, 05C65, 05D99
OR/MS subject classification: Primary: Networks/graphs, theory; secondary: mathematics, combinatorics
History: Received April 7, 2004; revised January 10, 2005.

1. Introduction. In a common situation in combinatorial optimization, there is a family of problems defined over the same set of possible solutions, where one wants an optimal solution for every problem in this family. Some of these cases capture an amount of degeneracy, that is, there are solutions that are optimal for more than one problem. In these kinds of *problems-solutions systems*, we raise the question of how big a set we need, in order to supply an optimal solution for every problem.

As an example, consider the multiterminal min-cut family of problems. Let $G=(V,E)$ be an edge-weighted graph with $|V|=n$. For $s,t\in V$, the (s,t) min-cut problem is to find a minimum weight cut separating s and t in G. Gomory and Hu [9] proved that there is a set of only $n-1$ cuts such that for each of the $\binom{n}{2}(s,t)$ problems there is at least one optimal solution in this set. We say that such a family of problems is $(n-1)$-*solvable*.

Hassin [11] and Cheng and Hu [4] proved a stronger version of this result. Consider a set X. Let a cut denote a partition of X into two sets. Assign to each cut an arbitrary weight. The problem is to find, for every pair $s,t\in X$ a minimum weight cut that separates it. Hassin [11] and Cheng and Hu [4] showed that even under this generalization, the problem is $(n-1)$-solvable.

An interesting generalization of this problem is the k-PAIRS problem, in which every problem consists of k pairs of elements, and the problem is to find an optimal cut that separates all the pairs simultaneously. Hassin [11] proved that the family of k-pairs problems on a set of n elements is d-solvable with $d=\sum_{m=1}^{k}\binom{n-1}{m}$. We note that for $k=1$ and $k=2$ the value of the minimum k-pairs cut equals the multicommodity max flow between the pairs, and for a bigger k it constitutes an upper bound on the flow.

This framework is not restricted to 0-1 problems. For instance, Hassin [11] considered the k-CUT problem, which is another generalization of the min-cut problem. In this problem, the input is a k-element subset of X, and a solution is a partition of the set into k nonempty subsets, each containing one element out of the input k-tuple. Hassin showed that this family of problems is d-solvable for $d=\binom{n-1}{k-1}$. Moreover, this d is the rank of the solution matrix for this problem (the solution matrix will be defined and discussed in the next section). Hassin generalized in this work a result of Lovász [16], giving this bound for the number of edges in a k-*forest*, which is a k-uniform hypergraph with the property that for each edge there is a k-cut separating all of its vertices and does not do so to any of the other edges. The edges of such a hypergraph can be viewed as a family of k-cut problems where each problem has a solution—a k-cut—that does not solve any other problem. Another proof for Lovász's result can be found in Parekh [17].

In some systems, we are able not only to prove the existence of the small-sized solutions set, but also to characterize its combinatorial structure, show an algorithm for finding it, and answer optimal solution queries using it. In the original work of Gomory and Hu [9], they also found a *compact representation* for the set of optimal solutions—a small-sized data structure containing all the $n-1$ solutions. They showed that if one builds a graph where each edge (u, v) is assigned the weight of the minimum cut between u and v, then there exists a spanning tree of this graph whose $n-1$ edges contain, for every problem (s, t), an edge corresponding to a minimum cut for (s, t). This highly regular structure is called the Gomory-Hu *cut tree*. Such a tree can be built by solving only $n-1$ min-cut problems. Given $s, t \in X$, the (s, t) min-cut can be computed out of the tree in $O(1)$ time if maintained properly.

Hassin [13] considered the family of *xcut* problems. In this set of problems one is given a pair s, t, and needs to find the optimal cut such that s and t are on the *same* side. Hassin showed a compact representation for all the optimal xcut solutions, analogous to the Gomory-Hu cut tree. This representation consists of a tree with one extra pair of elements, i.e., a base of the 2-forest matroid.

Hartvigsen [10] provided a compact representation for the k-cut and the k-pairs problems as well. This representation is based on matrices rather than graphs.

In the remainder of this section, we give an overview of the paper. In §2 we provide several definitions and notation that will be used throughout the paper. In §3 we give the necessary background for the three methods that we use to obtain the bounds. Section 4 contains our main results: In §4.1 we consider the family of k-*pairs* problems, all defined over a common ground set X of n elements. We first give a new proof for the result by Hassin [11] mentioned earlier, stating that there exists a set of size $\sum_{m=1}^{k}\binom{n-1}{m}$, which contains an optimal solution for every problem in this family. We then consider some subfamilies of this family of problems. For the family of *disjoint* pairs problems we show that a smaller set exists, one with $\binom{n}{k}-1$ solutions. Strengthening the latter result, we show that for the case where each problem has exactly k disjoint pairs (instead of at most k), a set containing an optimal solution to every problem exists with only $\binom{n}{k}-\binom{n}{k-1}$ cuts.

In §4.2, we look at the problem of 2-coloring a hypergraph with at most k edges. Again all problems considered are defined over the same set of n elements. This family is an extension of the k-pairs one. We show that although the number of problems in this family is much bigger than for the k-pairs problem, there exists a set containing an optimal solution to every problem of the same size as the one whose existence was proved for the k-pairs, namely $\sum_{m=1}^{k}\binom{n-1}{m}$. For the coloring problem of one hyperedge ($k=1$) of size at least t, we show the existence of an $(n-t+1)$-sized optimal solutions set, and that the Gomory-Hu tree of the complete graph on the n vertices contains such a set.

In §4.3, we prove the existence of a small optimal-solutions set for a family of problems we call the k-*in-one-side* problems. The problem here is again to find a minimum cut of the n elements, but this time we require that the cut puts k given elements on the same side. The size of the set whose existence we prove is $\sum_{m=0}^{k-1}\binom{n-1}{m}$. Next we look at the subfamily that contains instances with exactly k elements, and show that a smaller set exists for it—one with cardinality $\binom{n}{k-1}$. We derive from the latter bound a result by Lovász concerning critically chromatic hypergraphs, and show how to "translate" some of our other results into this language.

In §4.4, we consider the very general s-SAT problem. This is the problem of finding a satisfying assignment of minimum cost for a given Boolean formula in CNF with n variables and at most s clauses. In the price of finding a somewhat worse bound than the bounds in previous sections, we generalize a substantial amount of $\{0,1\}$ problems. For the s-SAT family of problems, we show the existence of a set of optimal solutions of size $\sum_{m=0}^{s}\binom{n}{m}$.

In §4.5, we investigate a special kind of problems-solutions systems that we call *take-out problems*, where one has to find the optimal subset of the ground set X after a number of elements have been removed from it, and possibly some other constraints have been applied. One such system is the MST-failure system, in which k edges may be taken out of a given graph with n vertices, and one has to find a minimum spanning tree for every instance. In this case we present a constructive combinatorial proof, i.e., one that tells us how to find the set of optimal solutions, whose size is $\binom{n-1+k}{n-1}$, and do so without solving unnecessary problems. We also apply a theorem from extremal combinatorics by Frankl [6], to show an extension of this bound for a wider collection of take-out systems, although the proof is nonconstructive.

2. Terminology. A *problems-solutions system* consists of a pair (P, S) as follows: S is the set of all potential solutions. This set is common to all problems. Typically, the solutions themselves are subsets or partitions of some ground set. A problem p is defined by the subset $S_p \subseteq S$ of solutions that are feasible for it—i.e., that conform with the problem constraints. P is a family of such problems, defined over the same set of solutions S. This defines a relation we call the *feasibility relation* of the system (P, S), which consists of the set of all pairs $\{(p_i, s_j)\}$ for which solution s_j is feasible for problem p_i, i.e., $s_j \in S_{p_i}$. Over the set of all possible solutions S, we define a *weight function* w: $S \to \mathbb{R}$ (or a cost function, depending on the context). We define w to be an arbitrary function, with no additional demands. In particular, it does not have to be nonnegative, additive, or one-to-one. Another way to look at w is as a weak total order, or a weak permutation over all solutions, determining for every two solutions whether one is better than the other. Solutions are allowed to have the same weight. The set of *optimal solutions* for problem p is defined by $S^*_{p,w} = \{s \mid s \in S_p \text{ and } \forall\, s' \in S_p\ w(s') \geq w(s)\}$. $S^*_{p,w}$ is the set of all feasible solutions of p with minimal weight (or sometimes maximal, again, depending on the context) among all other feasible solutions of p.

Given a problems-solutions system (P, S) and a weight function $w(S)$, a set of solutions $S' \subseteq S$ is called an *Optimal-Solutions Set* (or OSS for short) if for every $p \in P$ for which $S_p \neq \varnothing$, $S^*_{p,w} \cap S' \neq \varnothing$. An OSS is a set that contains at least one optimal solution for every problem in P, not including those problems in P for which there is no feasible solution at all. A minimal OSS is an OSS of minimum cardinality. While the latter is a property of the weight function, the next property we define is solely a property of the system (P, S): Let $W(S)$ be the set of all possible weight functions on S. The family of problems P is called *d-solvable* if for every $w \in W$ there exists an OSS S_w of cardinality $|S_w| \leq d$. The *solvability number* of the family of problems P is the minimum d for which P is d-solvable.

A hypergraph H is a collection of sets, called *edges*, whose members are called *vertices*. The *order* of a hypergraph is the number of vertices in the union of its edges. A hypergraph is called *k-uniform* if for each edge (i.e., set) $e \in H$, $|e| = k$. The *rank* of H is the maximum cardinality of a set in H. For a matrix A we will denote its rank over $\mathbb{F}_2$ by $\mathrm{rk}_2(A)$, and the rank over $\mathbb{R}$ by $\mathrm{rk}_{\mathbb{R}}(A)$. For a set of vectors T we will denote the subspace spanned by T by $\mathrm{sp}(T)$. The notation $\dot\cup$ stands for the disjoint union of sets. We will usually denote the ground set by X, and its size by n.

3. Methods. In this section, we develop some tools that will be useful in §4.

3.1. Triangular OSSs. Let (P, S) be a problems-solutions system.

Definition 3.1. We call an ordered set $T = \{(p_i, s_i)\}$ of problem and solution pairs *triangular* if:

$$\begin{cases} s_j \text{ is feasible for } p_i, & i = j; \\ s_j \text{ is not feasible for } p_i, & i > j. \end{cases} \tag{1}$$

PROPOSITION 3.1. *For every weight function w, there exists an OSS $\{s_i\}$ of (P, S) with different solutions weights, and a set of corresponding problems $\{p_i\}$, such that the set of pairs $T_w = \{(p_i, s_i)\}$, ordered according to w, is triangular.*

PROOF. We assume that the optimal solution has the lowest weight. For a given cost function $w \in W$, let $w(S)$ be the set of solution weights. For every $val \in w(S)$, let $S(val)$ be the set of solutions of weight val. Here is a greedy algorithm for finding T_w:

1. $T_w = \varnothing$.
2. For each val from $\min\{w(S)\}$ to $\max\{w(S)\}$ do:

2.1. While there exists a problem p_i, whose optimal solutions' weight is val, with no optimal solution in T_w, do:

2.1.1. Choose some optimal solution s_i of p_i and set $T_w = T_w \cup \{(p_i, s_i)\}$. $i = i + 1$.

The halting terms of Loop 2 and Loop 2.1 ensure that the set of solutions in T_w is indeed an OSS. To see that T_w is triangular, consider the step where the pair (p_i, s_i) was added to T_w. We want to show that for every $j < i$, s_j were not feasible for p_i. Indeed, if $w(s_j) = w(s_i)$, then by the term of Loop 2.1, s_j were not feasible for p_i. In addition, in the case $w(s_j) < w(s_i)$, if s_j would be feasible to p_i, then s_i could not have been an optimal solution of p_i. □

We will frequently use Proposition 3.1 via the following corollary:

COROLLARY 3.1. *For every cost function w, the minimum d for which P is d-solvable (i.e., the solvability number of P) equals the maximum cardinality of a triangular set in (P, S).*

PROOF. Let $t_{\max}$ denote the maximum size of a triangular problems-solutions pairs set. We have seen that for every weight function w, there exists a triangular OSS, thus the system is $t_{\max}$-solvable. To complete the proof, we need to show that for $t < t_{\max}$ the system is not t-solvable. Indeed, let T be the maximum-sized triangular set, and let w_T be a weight function giving the solutions in T the best (smallest) values. Then every OSS needs to contain all the solutions in T, so it has to be of size $t_{\max}$ or bigger, and the system is not t-solvable. □

The advantage in this point of view is that we no longer look at weight functions, but rather at a property of the feasibility relation of the system—the maximum cardinality of a triangular set in it.

Another two corollaries from Proposition 3.1 relate d-solvability to the number of different optimal solutions of a system:

COROLLARY 3.2. *If P is d-solvable, then there are no more than d distinct optimal solution values for the problems in P.*

COROLLARY 3.3. *If P is d-solvable and, in addition, every $w \in W$ is one-to-one, then the problems in P have no more than d different optimal solutions.*

3.2. The solution matrix. The solution matrix is a representation of the feasibility relation of the solutions to the problems.

DEFINITION 3.2. The solution matrix $A = A(P, S)$ of a system (P, S) is defined as follows:

$$A_{p,s} = \begin{cases} 1 & \text{solution } s \text{ is feasible for problem } p; \\ 0 & \text{otherwise.} \end{cases}$$

A square matrix is said to be *proper triangular* if it is triangular and its diagonal is all 1. By Corollary 3.1, finding the minimum d for which the family P is d-solvable can be reduced to the following formulation: What is the maximum size of a proper triangular submatrix over all permutations of the rows and columns of the solution matrix?

Finding the maximum proper triangular submatrix of a given matrix is NP-hard in general, as was shown in Bartholdi [3]. However, there exists a general bound on the *size* of the proper triangular submatrix: A proper triangular matrix is, in particular, a regular matrix, and this is true over any field $\mathbb{F}$. For a given matrix A, the maximum size of a regular submatrix (where the order of the rows and columns is not a factor) is nothing but the rank of A. Thus, the maximum size of a triangular submatrix of A is bounded by A's rank. In order to get the best bound, we will take the rank over $\mathbb{F}_2$, and we get the following useful theorem:

THEOREM 3.1 (HASSIN [11]). *Let (P, S) be a problems-solutions system. Then P is d-solvable for $d = \mathrm{rk}_2(A(P, S))$.*

In many of the systems we consider here and in Hassin [11], this bound is tight enough to find the solvability number of P. There are, however, instances in which using this bound is not such a good idea. For example, the matrix $J + I$ over $\mathbb{F}_2$, where J is the all 1s matrix and I is the identity matrix, is of almost full rank, but the maximum proper triangular submatrix of it is only of size 2.

We call a *solution basis* a basis of the row space of A. The solution basis is an interesting property of the system, for several reasons. First, it often gives the above-mentioned upper bound together with a lower bound on the size of the OSS. This is the case where the solution basis itself contains a triangular submatrix. Second, for a given weight function, one can look at the *maximum solution basis*, a basis whose total weight of solutions is maximal. From an algorithmic point of view, the maximum solution basis contains all the information needed to recover the solution to every problem. Hassin [12] gave the first algorithm for computing the maximum solution basis. This algorithm was later improved by Hartvigsen [10] to require solving only $\mathrm{rk}_2(A)$ number of problems. The solution basis often has a nice combinatorial structure, as is the case in the (s, t) min-cut system, with a general cut weight function. Hassin [11] showed that every solution basis for this system takes the form of a tree in the graph, and that problems associated with the edges of the Gomory-Hu tree form a maximum solution basis for these problems-solutions systems.

Finally, we define the *transposed system* of (P, S) to be (S, P), i.e., the problems become solutions and vice versa. It is easy to see that the solvability number of (P, S) equals the solvability number of (S, P): the solution matrices of two transposed systems satisfy $A(S, P) = A(P, S)^T$, and in addition, for every proper triangular matrix in $A(P, S)$, there is a proper triangular submatrix of the same size in $A(P, S)^T$ (when the order of the solutions is reversed from that of $A(P, S)$).

3.3. Spaces of polynomials. A useful method for bounding the number of objects in some configuration is associating a polynomial to each object. If we know the dimension of the space where these polynomials reside, then all we have to prove is their linear independence.

DEFINITION 3.3. Let (P, S) be a problems-solutions system. Let $\{Q_p \mid p \in P\}$ be a family of polynomials. We say that $\{Q_p\}$ are separating polynomials for the family P if there exists a set of vectors $\{v_s \mid s \in S\}$, such that $Q_p(v_s) = 0$ iff solution s is not feasible for problem p.

The following "triangular criterion" is a useful sufficient criterion for the independence of a set of functions.

PROPOSITION 3.2 (THE TRIANGULAR CRITERION, BABAI AND FRANKL [2]). *Let $S = \{f_i \mid f_i\colon \Omega \to \mathbb{F}\}$ be a set of functions from a set to a field, and let $e_i \in \Omega$ be elements such that*:

$$f_i(e_j) \begin{cases} \neq 0, & i = j; \\ = 0, & i > j. \end{cases}$$

Then the set S is linearly independent.

The way we will use this criterion is by associating a polynomial with each problem and a vector with each solution or the other way around. There is a special case where we can also bound the rank of the solution matrix using the dimension of the space where the polynomials lie:

LEMMA 3.1. *Let $\{Q_p\}$ be separating polynomials for the family of problems P, such that $\forall p\ Q_p \in \{0,1\}$. Let d be the dimension of* $\mathrm{sp}(\{Q_p\})$ *(taken either over $\mathbb{F}_2$ or over $\mathbb{R}$). Then* $\mathrm{rk}_2(A(P,S)) \le d$.

PROOF. The solution matrix here satisfies $A_{p,s} = Q_p(v_s)$, where v_s is the vector corresponding to solution s and Q_p is the polynomial of problem p. Note that $\dim_2 \mathrm{sp}(\{Q_p\}) \le \dim_{\mathbb{R}} \mathrm{sp}(\{Q_p\})$, so it is enough to show this lemma for the dimension taken over $\mathbb{F}_2$. Assume for contradiction that $\mathrm{rk}_2(A(P,S)) > d$. Then A contains a nonsingular submatrix B of size greater than d. The polynomials corresponding to the rows of B are linearly dependent and have some linear combinations equal to zero $\sum \lambda_p Q_p = 0$. Substituting the vectors a_j, we get a linear combination equals to zero, the same one for all columns. Thus, there is a linear dependence in the rows of B—a contradiction. □

Another relation between the solution matrix and the separating polynomials is the following:

PROPOSITION 3.3. *Let $\{Q_p\}$ be separating polynomials in n unknowns over $\mathbb{F}_2$ for the system (P,S). Assume that S is the power set of n elements. Let d be the dimension of* $\mathrm{sp}(\{Q_p\})$. *Then* $\mathrm{rk}_2(A(P,S)) = d$.

PROOF. We already know by Lemma 3.1 that $\mathrm{rk}_2(A(P,S)) \le d$ and would like to prove the other direction. Let $Q' \subseteq \{Q_p\}$ be a set of polynomials, such that $|Q'| > \mathrm{rk}_2(A(P,S))$. We want to prove that Q' is linearly dependent. Let $P' \subseteq P$ be the problems corresponding to Q'. The set of rows in the matrix corresponding to P' must be dependent, as it has more than $\mathrm{rk}_2(A(P,S))$ rows. Now take the linear combination to 0 of this set (which is, in fact, a sum) and apply it to Q'. Because S is the power set, this linear combination of Q' gives 0 $(\mathrm{mod}\, 2)$ for every point in the space; thus, it is the zero polynomial, and the polynomials are linearly dependent. □

Most of our results regard $\{0,1\}$ problems such as cuts, over a ground set X. In such problems the set of solutions S is the power set 2^X, so that Proposition 3.3 applies.

For a more comprehensive introduction to algebraic methods in combinatorics in general, and using polynomials in particular, see Babai and Frankl [2].

4. Bounds. Recall from §2 that for a system of problems-solutions (P,S), P is called d-solvable if for every weight function w there exists a set S_w of cardinality $|S_w| \le d$ containing an optimal solution to every problem (an OSS).

4.1. The k-pairs problem. Let X be a ground set of n elements. An instance of the k-PAIRS problem consists of $l \le k$ pairs of elements: $\{\{s_i,t_i\} \mid s_i \ne t_i\}_{i=1}^{l}$, $i = 1,\ldots,l$. A *feasible solution* to this problem is a cut $C = \{I_1, I_2\}$, where $X = I_1 \dot\cup I_2$, so that for every $1 \le i \le l$ either $s_i \in I_1$ and $t_i \in I_2$, or $s_i \in I_2$ and $t_i \in I_1$. For each cut C we associate a weight $w(C)$. The optimal solution to problem p is a feasible cut of minimum weight. We want to determine how big a set we need, in order to have a representative optimal solution to every k-pairs problem. A tight bound on the size of such a set is given by the following theorem:

THEOREM 4.1 (HASSIN [11]). *The family of k-pairs problems is d-solvable for*

$$d = \sum_{m=1}^{k} \binom{n-1}{m}. \tag{2}$$

Furthermore, d is the solvability number of this family.

Using different techniques, we present an alternative proof for the d-solvability of this family of problems, with the advantage of giving an improved bound for some interesting cases.

Proof. With each element $s_i \in X$, associate a binary variable x_{s_i}. With each problem p that consists of the pairs $\{\{s_i, t_i\}\}_{i=1}^{l}$, associate the following polynomial over $\mathbb{F}_2$: $Q_p = \prod_{i=1}^{l}(x_{s_i} + x_{t_i})$. Further, associate with each cut $C = \{I_1, I_2\}$ a vector v_C for which the jth component is 1 if the corresponding element is in I_1, and 0 otherwise (we call it the *incidence vector of* C).

Because the problems are symmetric in the sense that we do not distinguish the incidence vector of a solution and its complementary vector, we assume w.l.o.g. that $x_1 = 0$, so that we have $n-1$ unknowns. Clearly, all these polynomials reside in the space V of all polynomials of degree less than or equal to k over $\mathbb{F}_2$, whose dimension is $\sum_{m=1}^{k}\binom{n-1}{m}$. In addition, the polynomials $\{Q_p\}$ are separating for this family of problems, because Q_p vanishes by substituting an incidence vector of a cut C iff C is not a feasible solution for p. For a given cost function w, let $T_w = (P', S')$ be a triangular problems-solutions pairs set such that S' is an OSS. The existence of one is guaranteed in Proposition 3.1. By Proposition 3.2, the polynomials associated with the problems of T_w are linearly independent. Thus, there are at most $\sum_{m=1}^{k}\binom{n-1}{m}$ of them, and by Corollary 3.1 the theorem follows. □

Because the separating polynomials in the proof are over $\mathbb{F}_2$, by Lemma 3.1, the rank of the solution matrix for this system is also no more than $\sum_{m=1}^{k}\binom{n-1}{m}$.

The disjoint k-pairs problem is a special case of the k-pairs problem, in which no two pairs have a common element. An instance of this problem consists of $l \le k$ pairs: $\{s_i, t_i\}$ $i = 1, \ldots, l$, such that for every $i \ne j$, s_i, t_i, s_j, t_j are distinct elements. The feasibility relation is defined as before. We will assume $n \ge 2k$, so that it is meaningful to consider k disjoint pairs. For n elements, there are $\Theta(n^{2k})$ potential problems.

Theorem 4.2. *The family of disjoint k-pairs problems is d-solvable with*

$$d = \binom{n}{k} - 1. \tag{3}$$

Proof. For every cost function w, let $T_w = (P', S')$ be a triangular problems-solutions pairs set such that S' is an OSS. The existence of one is guaranteed in Proposition 3.1. Observe that because $n \ge 2k$, in every solution $s \in S'$, at least one of the sides of the cut contains at least k elements. Associate with every such solution a vector that is 1 for the elements of the bigger side and zeros for the other side. For a problem $p \in P'$ of separating the l pairs $\{\{s_i, t_i\}\}$, with a corresponding optimal solution in T_w, opt(p), consider first the polynomial $Q_p = \prod_{i=1}^{l}(x_{s_i} + x_{t_i})$, as in the proof of Theorem 4.1. Now, associate with the problem p the polynomial $Q'_p = Q_p \cdot x_{u_1}x_{u_2}\ldots x_{u_{k-l}}$, where $u_1, \ldots, u_{k-l}$ are arbitrary elements whose components in the vector associated with opt(p) are 1 and which do not appear in p. For problems of separating exactly k pairs, Q_p and Q'_p are the same. The polynomial Q'_p still vanishes under substitution of every solution that is not feasible for p, but not by substituting opt(p). (Note, however, that these polynomials are not separating for (P, S) anymore.) Thus, by the triangular criterion, these polynomials are linearly independent. The polynomials Q'_p lie in the space V of all polynomials of degree exactly k, and $\dim(V) = \binom{n}{k}$. Because we want to get a slightly smaller bound (by 1), we add the polynomial $Q_{\text{last}} = \prod_{i=1}^{k} x_i$ together with the all ones vector $\mathbf{1}$ at the beginning of the list of pairs. Obviously $Q_{\text{last}}(\mathbf{1}) \ne 0$, but $\forall p\ Q_p(\mathbf{1}) = 0$. Thus, all the polynomials together still satisfy the terms of the triangular criterion for independence of functions, and they are linearly independent in V.

It follows that there are no more than $\dim(V)$ polynomials altogether. Because for every problem p there exists such a polynomial, and because there exists one additional independent polynomial in V, we get the bound $\dim(V) - 1$ on the size of a triangular problems-solutions pairs set, and by Corollary 3.1 the theorem follows. □

We now turn to prove a stronger result, one that implies Theorem 4.2. However, the latter has the advantage of having a simpler proof. The DISJOINT EXACT k-PAIRS problem is a further restriction of the k-pairs problem. An instance of this problem is composed of *exactly* k-pairs. The feasibility relation remains as before, namely, a cut solves the problem iff it separates every pair in it. We still assume $n \geq 2k$. To prove d-solvability for this family of problems, we need a lemma. The following definition is a variation of one that appears in Jukna [14]:

DEFINITION 4.1. Let k, l, n be three natural numbers such that $k, l \leq n$, and let X be a set of n elements. The (n, k, l) *disjointness matrix* $D_n(k, l)$ over X is a 0-1 matrix whose rows are labeled by subsets of X of size *exactly* k and whose columns are labeled by subsets of X of size *at most* l. The entry $D_{A,B}$ in the Ath row and Bth column is defined by:

$$D(A, B) = \begin{cases} 0, & A \cap B \neq \varnothing; \\ 1, & A \cap B = \varnothing. \end{cases}$$

LEMMA 4.1. *Suppose $k + l \leq n$, and let $m = \min(k, l)$. Then,* $\mathrm{rk}_2 D_n(k, l) = \binom{n}{m}$.

We discuss and prove Lemma 4.1 in Appendix A.

THEOREM 4.3. *The family of disjoint exact k-pairs problems is d-solvable with*

$$d = \binom{n}{k} - \binom{n}{k-1}. \tag{4}$$

PROOF. Associate with each problem the polynomial $Q_p = \prod_{i=1}^{k}(x_{s_i} + x_{t_i})$. As mentioned in the proof of Theorem 4.1, the polynomials associated with a triangular OSS of this family are linearly independent, and the dimension of the polynomial space in which they lie is easily seen to be $\binom{n}{k}$. However, we will show that there are additional $\binom{n}{k-1}$ polynomials, such that the whole set is linearly independent in that space. We will add all monic monomials $\{Q'_i\}$ of degree exactly k. Consider the following matrix, which is composed of four submatrices $B = \begin{pmatrix} B_1 & B_2 \\ B_3 & B_4 \end{pmatrix}$. The rows of B_1 and B_2 are labeled by the polynomials Q_p. The rows of B_3 and B_4 are labeled by the polynomials Q'_i. The columns of B_2 and B_4 are labeled by binary vectors with strictly less than k zeros, and the columns of B_1 and B_3 are labeled by vectors with k zeros or more. The item $B_{i,j}$ is the result of substituting the vector of the column j in the polynomial of the row i. First note that B_2 is an all-zeros matrix, because every Q_i, being a separating polynomial for an exact k-pair problem, vanishes by substituting a vector with less than k zeros.

Thus, we get that $\mathrm{rk}_2(B_1) + \mathrm{rk}_2(B_4) \leq \mathrm{rk}_2(B)$. By Lemma 3.1, because all polynomials reside in a space of dimension $\binom{n}{k}$, $\mathrm{rk}_2(B) \leq \binom{n}{k}$. Thus, $\mathrm{rk}_2(B_1) + \mathrm{rk}_2(B_4) \leq \binom{n}{k}$. We will next show that $\mathrm{rk}_2(B_4) = \binom{n}{k-1}$. It is easy to see that the entry $B_{4_{i,j}}$ is 0 iff the k-tuple of the monomial Q'_i intersects the $<k$ zeros of the vector a_j. Thus, B_4 is an $(n, k, k-1)$ disjointness matrix, so by Lemma 4.1, $\mathrm{rk}_2(B_4) = \binom{n}{k-1}$. Because B_2 is all zeros, it follows that $\mathrm{rk}_2(B_1 \mid B_2) = \mathrm{rk}_2(B_1) \leq \binom{n}{k} - \binom{n}{k-1}$. Notice that the matrix $(B_1 \mid B_2)$ is nothing but the solution matrix for this system, so by Theorem 3.1, this family is d-solvable for $d = \binom{n}{k} - \binom{n}{k-1}$. □

4.2. 2-Coloring a hypergraph with at most k edges. The problem of 2-coloring a hypergraph has been considered in various contexts in combinatorics. In its most common form, a problem instance is a hypergraph H, and a feasible solution is a coloring of its nodes in two colors, where no edge is monochromatic. For each coloring C out of the 2^{n-1} possible colorings (we do not distinguish opposite colorings) we associate a weight $w(C)$. The problem is to find a feasible coloring of minimum weight. Note that this problem generalizes the k-pairs problem, in which every edge is of size 2. We are interested in bounding the number d for which this family of problems is d-solvable.

THEOREM 4.4. *The family of hypergraph 2-coloring problems on hypergraphs with at most k edges is d-solvable for*

$$d = \sum_{m=1}^{k} \binom{n-1}{m}. \tag{5}$$

Furthermore, d is the solvability number of this family.

PROOF. We prove the theorem by a reduction to the k-pairs problem.[1] Note that the two problem families—the coloring and the k-pairs—are defined over the same set of solutions. We show that each optimal solution for the coloring problem is also an optimal solution for some instance of the k-pairs problem. Let H be a hypergraph with l edges ($l \le k$), p its coloring problem, and C an optimal coloring for it. Let $p' = \{\{s_i, t_i\}\}_{i=1}^{l}$ be l pairs, not necessarily distinct; each pair belongs to the same edge, which are separated by this coloring (an arbitrary choice out of all possible ones). For a given cost function w, denote the set of optimal cuts for p' by $S^*_{p',w}$. *Claim*: Every optimal solution for p' is also an optimal solution for p. *Proof*: *Fact* 1: Every solution C' for the k-pairs problem p' is also a feasible solution for the coloring problem p, in particular, $S^*_{p',w}$ is feasible for p. *Fact* 2: Every $s \in S^*_{p',w}$ has the lowest weight among all solutions feasible for p. This is true because there exists (by the way p' was chosen) at least one optimal solution C for p that is feasible for p', so the optimal solution weight for p' must be smaller than or equal to it. From these two facts together, the claim follows.

Let O be an OSS for the k-pairs problems under w. We have shown that for each coloring problem p there exists a k-pairs problem p' such that the set of optimal solutions for these two problems satisfies $S^*_{p',w} \subseteq S^*_{p,w}$. Because O contains an optimal solution for p' out of $S^*_{p',w}$, it contains one also for p. Thus, O is an OSS for the coloring problems as well. Because by Theorem 4.1 the k-pairs family of problems is d-solvable for $d = \sum_{m=1}^{k} \binom{n-1}{m}$, the same is true for the coloring family of problems. □

The result of Theorem 4.4 is tight, as it generalizes Theorem 4.1.

Note that if all the hypergraphs considered have disjoint edges, then by a reduction to the disjoint k-pairs family of problems, we get the improved bound of $d = \binom{n}{k} - 1$ (without the "tail"). If in addition these hypergraphs have exactly k edges, then by reduction to the exact-disjoint k-pairs problem, we get $d = \binom{n}{k} - \binom{n}{k-1}$.

A generalization of Theorem 4.4 can be made such that the term "H has at most k edges" is replaced by a weaker one: For a hypergraph H let the *inclusion number* of H be the minimum number of hyperedges in a subhypergraph H' of H, in which for every edge $e \in H$ there exists $e' \in H'$ such that $e' \subseteq e$.

PROPOSITION 4.1. *The family of hypergraph 2-coloring on hypergraphs with inclusion number at most l is d-solvable for $d = \sum_{m=1}^{l} \binom{n-1}{m}$.*

PROOF. The proof is the same as that of Theorem 4.4, except that we take the same pair $\{s_i, t_i\}$ for all edges containing the same edge e' in H' (if two edges contain the same e', then we choose one arbitrarily). These pairs constitute p', which is a k-pairs problem with $k = l$. □

Let us also consider in more detail the case $k = 1$ for the hypergraph 2-coloring problem (which might also be called THE CUTTING EDGE problem . . .): A problem instance now is a single hyperedge e (a set), and a feasible solution to it is a coloring of the n elements in two colors, such that e is not monochromatic. First, a nice result of the proof of Theorem 4.4 is based on the fact that for $k = 1$ each optimal solution to this problem is an optimal

[1] In the special case where the hypergraphs are q-uniform, where q is prime, we can use the same technique as in the k-pairs problem, by considering the polynomials $\prod_{e\in H} \sum_{i\in e} x_i$ over $\mathbb{Z}_q$. For q nonprime (or prime power) $\mathbb{Z}_q$ is not a field. The k-pairs problem is the case $q = 2$.

solution for some (s, t) min-cut problem. Thus, given the set of elements X and the weight function w, the Gomory-Hu cut tree for all the (s, t) min-cut problems on X forms an OSS for the family of problems we consider.

In this case of one hyperedge, if every problem instance is of size at least t, then one can prove an even better bound that improves as a function of t:

PROPOSITION 4.2. *The family of* 2*-coloring a single hyperedge of size at least t problems is* $(n-t+1)$*-solvable.*

PROOF. For a given cost function w, let $T_w = \{(p_i, s_i)\}$ be a triangular problems-solutions pairs set such that $S' = \{s_i\}$ is an OSS. The existence of one is guaranteed in Proposition 3.1. Assume that T_w is ordered in the opposite way than the one found in Definition 3.1 (so that s_1 is the solution of highest weight in S'—i.e., the worst one). For the t-sized hyperedge problem p_1, all solutions in $S' \setminus \{s_1\}$ are not feasible, by the triangularity of T_w. Therefore, for every cut in $S' \setminus \{s_1\}$ all elements of p_1 form a cluster of elements that belong to the same side. Similarly, for every cut in $S' \setminus \{s_1, s_2\}$ all elements of p_1 form a cluster, and all elements of p_2 form another cluster. In each step other than the first, a t-tuple of elements that were not previously all in the same cluster stick together to form one cluster for the next steps, i.e., at least two former clusters become one, so the number of clusters reduces by one. In the first step t elements become one cluster, so the number of clusters reduces by $t-1$. Thus, there cannot be more than $n-t+1$ steps and the proposition follows. □

This simple proof still generalizes, by taking $t = 2$, the case considered by Cheng and Hu in [4] and by Hassin in [11], that is, the extension of the Gomory-Hu bound for arbitrary cut weights.

4.3. The k-in-one-side (k-xcut) problem. Let $C = \{I_1, I_2\}$, where $X = I_1 \dot{\cup} I_2$, be a cut. We say that a set $R \subseteq X$ *crosses* C if $R \cap I_1 \neq \varnothing$ and $R \cap I_2 \neq \varnothing$. The k-IN-ONE-SIDE problem is a generalization of the xcut problem posed by Hassin [13], in which $k = 2$. Given a universe X of n elements, an instance of this problem is a set $R \subset X$, such that $|R| \leq k$. A feasible solution to problem p, represented by the set R, is a cut $C = \{I_1, I_2\}$ where $X = I_1 \dot{\cup} I_2$, such that R does not cross this cut, i.e., either $R \subseteq I_1$ or $R \subseteq I_2$. Given a weight function w over the cuts, the problem is to find a feasible cut for R of minimum weight. There are $\sum_{m=1}^{k} \binom{n}{m}$ k-in-one-side problems.

THEOREM 4.5. *The family of k-in-one-side problems is d-solvable for*

$$d = \sum_{m=0}^{k-1} \binom{n-1}{m}. \tag{6}$$

PROOF. Associate with each element $e_i \in X$ a binary variable x_i and associate with each problem p the following polynomial over $\mathbb{F}_2$: $Q_p = \prod_{e_i \in p} x_i + \prod_{e_i \in p} (1 + x_i)$. With each solution s, associate as always the incidence vector of its cut v_s. Obviously Q_p vanishes by substituting the incidence vector of solutions that are not feasible for p, and only by such solutions. The term with highest degree, $\prod_{e_i \in p} x_i$, appears exactly twice in the sum, so the polynomial over $\mathbb{F}_2$ is of degree $k-1$. Because the problems are symmetric, in the sense that we do not distinguish the incidence vector of a solution and its complementary vector, we may assume w.l.o.g that $x_1 = 0$. Thus, we have only $n-1$ unknowns. We get that the dimension of the space here is $d = \sum_{m=0}^{k-1} \binom{n-1}{m}$. For a given cost function w, let $T_w = (P', S')$ be a triangular problems-solutions pairs set such that S' is an OSS. Then by Proposition 3.2, the polynomials corresponding to problems in P' are linearly independent, so there are no more than d of them, and $|P'| \leq d$. □

Because the polynomials are separating over $\mathbb{F}_2$, by Lemma 3.1 we also get that the rank of the solution matrix for this system is bounded by $\sum_{m=0}^{k-1} \binom{n-1}{m}$. In the following theorem we will also see that this bound is tight.

Recall that a solution basis for a problems-solutions system is a basis for the row space of its solution matrix. We would now like to find a solution basis for this family of problems. This result extends Theorem 4.5 as it also gives another proof for the d-solvability of this family of problems.

THEOREM 4.6. *Let $1 \in X$ be an arbitrary element, let P' be the family of all k-in-one-side problems defined by the sets $\{R \mid R \subseteq X,\ |R| \le k,\ 1 \in R\}$, and let $B_{P'}$ be the corresponding rows in the solution matrix $A(P, S)$ of this system. Then $B_{P'}$ constitutes a solution basis for the family of all k-in-one-side problems. Thus, $\mathrm{rk}_2(A(P,S)) = d = \sum_{m=0}^{k-1}\binom{n-1}{m}$. Furthermore, d is the solvability number of the k-in-one-side family.*

PROOF.
Claim 1. The rows of $B_{P'}$ are linearly independent over $\mathbb{F}_2$. Proof: For each problem p composed of the set $R = (1, i_1, \ldots, i_{k'}) \in B$ where $0 < k' < k$, the cut $C = (R, X \backslash R)$ solves it, but no other problem $p' \in P'$ satisfying $|p'| \ge |p|$.
Claim 2. The rows of $B_{P'}$ span the rows of the solution matrix. Proof: Let p be a problem not in P', i.e., $R = \{j_1, \ldots, j_{k'}\}$, where $1 \le k' \le k$ and $1 \notin R$. Consider the set of problems $P_0 = \{\{1\} \cup R' \mid R' \subset R\}$, and its corresponding set of rows B_0. Note that $B_0 \subseteq B$ and that $|B_0| = 2^{k'} - 1$. We will show that the rows corresponding to B_0 span the row of p. For the claim to hold, we show that each cut solves an even number of problems in $P_0 \cup \{p\}$. Consider first a cut $C = \{I_1, I_2\}$, where $X = I_1 \dot{\cup} I_2$, that solves p. Assume w.l.o.g that $R \subseteq I_1$ (the other option being $R \subseteq I_2$). If $1 \notin I_1$ then C solves only p and $\{1\}$—an even number of problems. Otherwise, if $1 \in I_1$, then C solves all the $2^{k'} - 1$ problems in P_0 plus p—again, an even number of problems. Consider next a cut C that does not solve p, i.e., p crosses C. Assume w.l.o.g that $1 \in I_1$ and let $R = p \cap I_1$ and $|R| = k''$ where $0 < k'' < k'$. C solves only the problems $\{\{1\} \cup R' \mid R' \subseteq R\}$. Thus, C solves exactly $2^{k''}$ problems—again, an even number. To summarize, no matter if C solves p or not, it always solves $0 \pmod 2$ problems from $B_0 \cup \{p\}$. Thus, the rows corresponding to $B_0 \cup \{p\}$ are dependent, and because the rows of B_0 are independent, we get that B_0 spans p. Thus, B spans p, and we get that B forms a solution basis.
By the proof of Claim 1, the bound of Theorem 4.5 is tight, for the problems corresponding to this solution basis, together with their optimal solutions, form a triangular set of size $\sum_{m=0}^{k-1}\binom{n-1}{m}$. Thus, this is the solvability number of this family. □
Consider now the set of *uniform problems*, for which $|R| = k$ (instead of $|R| \le k$). We can get a somewhat better bound for the d-solvability of these problems only.

THEOREM 4.7. *The family of uniform k-in-one-side problems on a set of n elements is d-solvable for*

$$d = \binom{n}{k-1}. \tag{7}$$

The proof is based on ideas from Alon et al. [1].
PROOF. This time we associate polynomials with solutions. Associate with each element e_i the variable x_i. With each problem $p = (e_1, \ldots, e_k)$ we associate its incidence vector x_p. With each solution we associate its incidence vector v_s and the following polynomial over $\mathbb{R}$: $Q_s = \prod_{t=1}^{k-1}(v_s x - t)$. Q_s vanishes by substituting x_p iff v_s is the incidence vector of a cut that does not solve p. For a given cost function w, let $T_w = (P', S')$ be a triangular problems-solutions pairs set such that S' is an OSS for this family of problems, ordered opposite to the order of Definition 3.1. Then, by the triangular criterion for independence of functions, the polynomials associated with all $s \in S'$ are linearly independent. The polynomials reside in a space of dimension $\sum_{m=0}^{k-1}\binom{n}{m}$. To get the smaller bound, we now add to this list of polynomials/vectors pairs, several other pairs at the end of the list: For every subset I of the ground set satisfying $|I| \le k-2$, we add the polynomial $\prod_{i\in I} x_i(\sum_{j=1}^{n} x_j - k)$, along with the incidence vectors of I, where the pairs are ordered by the cardinality of I

from smallest to biggest. The whole set of polynomials/vectors remains triangular and thus, by the triangular criterion, the polynomials are linearly independent. Therefore, the number of polynomials before the addition must not be greater than $\binom{n}{k-1}$. □

Whether the number in Theorem 4.7 is indeed the solvability number of the system remains an open problem.

Note that neither of the Theorems 4.5 and 4.7 follow from the other. The relation between the bounds is $\sum_{m=0}^{k-1}\binom{n-1}{m} = \binom{n}{k-1} + \sum_{m=0}^{k-3}\binom{n-1}{m}$. Note also that Theorem 4.7 does not imply a bound on the rank of the solution matrix for this system.

For the case $k = 2$, the xcut problem, Hassin [13] proved that the system is d-solvable for $d = n$. This result is easily seen to be a special case both for Theorem 4.5 and Theorem 4.7. This is the one case where these two bounds coincide.

From Theorem 4.7 we can derive a result by Lovász [16]: The *chromatic number* of a (hyper)graph H is the minimum number of colors needed to color the vertices of H in a legal way, i.e., in a way for which no hyperedge is monochromatic. We call H *critically l-chromatic* if it is l-chromatic, but removing any edge from H yields a (hyper)graph that can be colored properly with fewer than l colors.

Corollary 4.1. *Let H be a critically 3-chromatic k-uniform hypergraph of order n with edge set E. Then $|E| \le \binom{n}{k-1}$.*

Proof. Removing any edge creates a 2-chromatic hypergraph. Thus, for each edge there is a cut that crosses all edges but it. In particular, looking at each edge as a k-in-one-side problem, and assigning each cut with a different weight w, these problems have different optimal solutions. Thus, an OSS for this weight function has to contain a solution for every problem—that is, for every edge. By Theorem 4.7, a family of k-in-one-side problems is $\binom{n}{k-1}$-solvable, thus $|E| \le \binom{n}{k-1}$. □

Note that in Theorem 4.7 we get the bound under weaker terms than those assumed in Lovász [16]: Lovász's result demands that the hypergraph H is not 2-colorable, that is, there is no cut separating all edges, whereas ours does not, as well as the skew (triangular) nature of our theorem, which is not obtained by any easy modification of Lovász's proof.

Now, we can derive an analogous result from Theorem 4.5:

Corollary 4.2. *Let H be a critically 3-chromatic hypergraph of order n and rank k, with edge set E. Then $|E| \le \sum_{m=0}^{k-1}\binom{n-1}{m}$.*

Continuing this line, we translate Proposition 4.2 to these terms as well:

Corollary 4.3. *Let H be a t-uniform hypergraph of order n with edge set E. If for every $e \in H$ there exists a coloring C in which e is the only nonmonochromatic edge, then $|E| \le n - t + 1$.*

To end this section, consider the problem of finding the minimum weight *illegal* 2-coloring for a hypergraph of order n and rank k, which we name the noncoloring problem.

Proposition 4.3. *The noncoloring problem is d-solvable for $d = \sum_{m=0}^{k-1}\binom{n-1}{m}$.*

Proof. The proof is by reduction to the k-in-one-side problem. Let H be a hypergraph of rank k, p a noncoloring problem on H, and C an optimal illegal coloring for it. Let e be one of the monochromatic edges of H in C. For a given cost function w, consider the set of optimal cuts for the uniform k-in-one-side problem defined by e. It is easy to see that every optimal solution for e is also an optimal solution for p, and the proposition follows. □

An asymmetric variant of the noncoloring problem is the set inclusion problem. A problem instance is a k-uniform hypergraph H, and a feasible solution is a set of minimum weight that contains at least one of H's edges.

Proposition 4.4. *The family of all set inclusion problems is d-solvable, with $d = \binom{n}{k}$.*

Proof. Associate with every element $i \in X$ a variable x_i, and with every hypergraph H the polynomial: $Q_H = \sum_{e\in H} \prod_{i\in e} x_i$. Associate further with every solution its incidence vector. $\{Q_H\}$ are easily seen to be separating polynomials for this family of problems, and they lie in a space of dimension $\binom{n}{k}$. □

We note that a proof using reduction could be applied here as well.

4.4. SAT problems. A very general $\{0, 1\}$ problem is the one of satisfying a Boolean formula. We call a Boolean function *s-CNF-able* if it can be written in CNF form with at most s clauses.

For a given s-CNF-able formula f on n variables, and a weight function w: $2^n \to \mathbb{R}$ that associates with each truth assignment of these variables an arbitrary weight, the s-SAT problem is that of finding a minimum weight satisfying assignment to f.

Theorem 4.8. *The family of s-SAT problems is d-solvable for*

$$d = \sum_{m=0}^{s} \binom{n}{m}. \tag{8}$$

Furthermore, d is the solvability number of this system.

Proof. Denote the rth clause in an arbitrary problem p by C_r. With p associate the polynomial $Q_p = \prod_{r=1}^{s} (\sum_{x_i\in C_r} x_i + \sum_{\bar{x}_i\in C_r} (1 - x_i))$ over $\mathbb{R}$. Each such polynomial vanishes by substituting x for an incidence vector of a solution not feasible to p, and does not vanish by substituting a feasible solution. For a given weight function w, let $T_w = (P', S')$ be a triangular problems-solutions pairs set such that S' is an OSS. The existence of one is guaranteed in Proposition 3.1. Then, by Corollary 3.1, the terms of the theorem and the triangular criterion, polynomials associated with problems in P' are linearly independent. In addition, all these problems lie in a space of dimension $\sum_{m=0}^{s} \binom{n}{m}$ and the upper bound follows.

This bound is tight, as can be seen from taking the conjunction of every set of $m \le s$ literals with no negation, out of the given n. For every conjunction, the assignment giving all its literals the value "true," and all the rest the value "false," satisfies it, but no other formula with fewer or the same number of literals. For $m = 0$, we take a formula with negation only (or alternatively, an empty formula), together with the all-false assignment. □

Following Jukna [14], a Boolean function f in DNF form is called *s-or-and* function if every clause in it has at most s literals. There is an immediate corollary regarding the optimal assignment problems for such functions over a set of n variables:

Corollary 4.4. *The family of s-or-and problems is d-solvable for*:

$$d = \sum_{m=0}^{s} \binom{n}{m}. \tag{9}$$

Proof. By applying the distributive law, every s-or-and function can be written as a CNF formula with at most s clauses and vice versa. Thus, they are in fact the same family of problems, and by Theorem 4.8 the s-SAT family is $\sum_{m=0}^{s} \binom{n}{m}$-solvable. □

Another result one can draw from Theorem 4.8 and its corollary regards a family of formulas in propositional calculus where there are at most r literals appearing in every formula. We call such formulas *r-literals* formulas. Note that for this corollary to hold, one does not have to assume a certain form of the formula:

Corollary 4.5. *The family of r-literals formulas is d-solvable with*

$$d = \sum_{m=0}^{r} \binom{n}{m}. \tag{10}$$

PROOF. For every r-literals formula, its DNF form has at most r literals in every clause. Now apply Corollary 4.4. □

Our general results on SAT problems can be given some interesting interpretations. Here is one example: A *transversal* in a hypergraph, sometimes also called a *blocking set*, is a set that intersects every edge in the hypergraph. The s-SET-COVER problem obtains as an input a hypergraph H with at most s edges and outputs a set of minimum weight which is a transversal of H. Note that every instance of this problem can be seen as an s-SAT problem, whose literals appear with no negation (such formulas are also called *monotone*), so we have the following interpretation:

PROPOSITION 4.5. *The s-set-cover family of problems on hypergraphs with at most k edges is d-solvable with $d = \sum_{m=1}^{k} \binom{n}{m}$. Furthermore, this is the solvability number of this family.*

The proof is similar to that of Theorem 4.8.

REMARK. A transversal is a set that does not have an intersection of size 0 with any edge in H. Proposition 4.5 can be easily extended to the problem of finding a set that avoids intersections of size u.

4.5. Take-out problems. Let $G = (V, E)$ be a multigraph. Suppose that a failure may occur for at most k edges in G. We would like to be prepared for these failures and find a minimum spanning tree (MST) for all possible resulting connected graphs in advance. How many minimum spanning trees must we find and keep? We call this the MST-FAILURE family of problems.

For every such multigraph $G = (V, E)$ and edges $E' \subseteq E$, we define $G' = (V, E \backslash E')$ to be the graph after the failure of these edges. For simplicity we assume $k < n$.

LEMMA 4.2. *Let $T = (V_T, E_T)$ be an MST of G. Let E' be the set of failed edges, and denote by $E_0 = E_T \backslash E'$ the remaining edges of T. Then there exists an MST of G' that contains E_0.*

PROOF. Because E_0 are edges of an MST of G, there is no cycle of G in which one of these edges is strictly the most expensive one (otherwise, there was a cheaper MST to G than the one assumed). Let T' be some MST of G'. We iteratively add each edge in $E_0 \backslash E_T$ to T'. Every such edge e closes a cycle in T', in which there is at least one edge not in E_0 that is the most expensive one. Deleting the expensive edge, we again get an MST. After the last iteration, we get an MST of G' containing all the edges in E_0. □

THEOREM 4.9. *The MST-failure family of problems has solvability number $d = \binom{n-1+k}{n-1}$. Furthermore, the OSS of a given weight function w can be constructed by solving only d problems.*

PROOF. By induction on k, the number of edges removed. For $k = 0$ the claim is trivial. Now, assume the claim is true for all values smaller than k and we took out at most k edges: i out of the MST of G and at most $k - i$ other edges. For $i = 0$, the MST remains as is. Now consider the case $i > 0$, and let E_0 be the MST edges not removed. By Lemma 4.2, the edges of E_0 remain tree edges together in some new tree of G'. Thus we may shrink every connectivity component of the tree not removed into a single vertex. Because there exists an MST of G' that contains all the edges of E_0, obviously such a tree can be found by connecting these components in the optimal way. We get a new multigraph G'' with $i + 1$ vertices. By the induction hypothesis, there is a set containing an MST for every graph that may be created out of G'' by taking out the remaining $k - i$ edges, and the cardinality of this set is at most $\binom{k-i+i}{k-i} = \binom{k}{i}$. Summing over all values of i, the cardinality of the set containing all MSTs obtained from G by removing at most k edges is $\sum_{i=0}^{k} \binom{n-1}{i}\binom{k}{i} = \binom{n-1+k}{k}$.

To see that this bound is tight, consider a multigraph on n vertices numbered $1, \ldots, n$ where for every $1 \le i \le n-1$, there are $k+1$ differently weighted edges connecting vertex i to vertex $i+1$, and these are the only edges. Consider a vector $J = (j_1, \ldots, j_n)$ such that $\sum_{i=1}^{n} j_i = k$. Taking out the j_i lightest edges between vertex i and $i+1$ creates some new MST. For every choice J we get a different MST. j_n completes the sum to be exactly k. Because $\sum_i j_i = k$, the number of possibilities of choosing the numbers j_i is $\binom{n-1+k}{k}$. □

A generalization of the MST-failure problem can be made as follows: Let X be some ground set of elements, with a weight function defined on its power set. A problem instance consists of a set A of forbidden elements, together with the corresponding collection of allowed subsets $\mathcal{F} = 2^{X \setminus A}$, some of which may have infinite weight. The task is to find a subset $B \in \mathcal{F}$ of minimum weight. We call such problems TAKE-OUT problems. A bound may be given for a very general family of take-out problems using a theorem from extremal combinatorics:

THEOREM 4.10 (FRANKL [6]). *Let $A_1, \ldots, A_m$ and $B_1, \ldots, B_m$ be finite sets such that $A_i \cap B_i = \varnothing$ and $A_i \cap B_j \neq \varnothing$ for $j < i$. Suppose also that $|A_i| \le k$ and $|B_i| \le l$. Then $m \le \binom{k+l}{k}$.*

The consequence for take-out problems is immediate:

COROLLARY 4.6. *Let (P, S) be a system of take-out problems. Suppose that $\forall\, i\ |P_i| \le k$ and $\forall\, i\ |S_i| \le l$. Then the family P is d-solvable for $d = \binom{k+l}{k}$.*

This result generalizes the bound of Theorem 4.9. However, the proof of Theorem 4.10 (as is the case with many of the above bounds) is nonconstructive, i.e., it does not tell us how to find the set of all the optimal solutions whose existence it guarantees. We note also that there is no combinatorial proof known for the general case. In the special case of MSTs, we showed a constructive simple combinatorial proof for this bound. In addition, we showed that although the terms of the problem were stricter, the result remained tight. An analogous result regarding take-out shortest paths can be found in Eilam-Tzoreff [5]. Given two vertices, the problem there is to find all shortest paths between them after at most k edges have been removed.

The family of take-out problems considered in Corollary 4.6, with $k = l$, is a natural example for a case where the bound of Theorem 3.1 gives a very poor result: The solution matrix of such a system has full row rank, that is, $\sum_{i=0}^{k} \binom{n}{i}$ (see Razborov [18]), whereas this family is $\binom{2k}{k}$-solvable (from Theorem 4.10), no matter how big the size n of the ground set is.

Appendix A. Proof of Lemma 4.1. Lemma 4.1 can also be derived from other more general results in the literature (see e.g., Frankl [7], Frumkin and Yakir [8], Linial and Rothschild [15], and Wilson [19]). We derive the result here in a different way, which we believe is a bit simpler. Recall that the (n, k, l) disjointness matrix $D_n(k, l)$ over X is a 0-1 matrix whose rows are labeled by subsets of X of size *exactly* k, whose columns are labeled by subsets of X of size *at most* l, and whose entry $D_{A,B}$ is 1 iff A and B are disjoint.

LEMMA A.1. $\mathrm{rk}_2\, D_n(k, k) = \binom{n}{k}$.

PROOF (RAZBOROV [18] AND JUKNA [14]).[2] We will show that the rows of D are linearly independent. Any nontrivial linear combination of the rows of D is a sum over $\mathbb{F}_2$ of a subset of the rows. Let M be the nonempty set of k-tuples associated with these rows. Now consider the polynomial $q = \sum_{I \in M} \prod_{i \in I} x_i$. Take some arbitrary k-tuple $I_0 \in M$ and substitute in q, $x_i = 1$ for every $i \notin I_0$ to get a new polynomial q'. Because the term $\prod_{i \in I_0} x_i$

[2] Actually Razborov has proved a stronger result, stating that the matrix of at most k-tuples against at most k-tuples has full row rank.

is left untouched, q' is a nonzero polynomial. Thus, there exists an assignment b' for the k variables of q', for which $q'(b')=1$. Let b be the extension of that assignment to all other variables, so we know $q(b)=1$. Let J_0 be the set of elements whose variables are assigned zero in b. Because $q(b)=1$, the number of nonzero additives in the sum $q(b)$ is odd. However, this is the number of k-tuples disjoint to J_0. We get that for every subset of the rows of D, there exists a column (labeled by J_0) for which the sum of rows is $1 \pmod 2$. Thus, every nontrivial linear combination of D's rows is nonzero. □

Lemma A.2. *Let k, l, n be natural numbers such that $k+l\le n$, and let $m=\min(k,l)$. Then* $\mathrm{rk}_2 D_n(k,l)\ge\binom{n}{m}$.

Proof. We apply induction on k and l. The base cases are $k=0$ and $l=0$, and in these cases the matrix is a single row/column of ones, so the rank is trivially 1. Another case we already know to be true is the case $k=l$, for which we have Lemma A.1, so in the following we will prove the lemma only for $l\ne k$. Our induction hypothesis is that for every $k'\le k$ and $l'<l$, and for every $k'<k$ and $l'\le l$, if $n\ge k'+l'$, then $\mathrm{rk}_2(D(k',l',n))\ge\binom{n}{m'}$ for $m'=\min(k',l')$. Consider now the matrix $D_n(k,l)$ and its submatrix D_1 that consists of rows corresponding to all subsets of size at most k that include some specific element e, and columns corresponding to subsets of size exactly l that *do not* include e. D_1 is an $(n-1,k,l-1)$ disjointness matrix over $X\backslash\{e\}$, which satisfies the terms of the lemma. Thus, by the induction hypothesis, $\mathrm{rk}_2(D_1)\ge\binom{n-1}{\min(k,\ l-1)}$. Now consider the submatrix D_2, whose rows correspond to all subsets of size at most k that do not include e and whose columns correspond to subsets of size exactly l that include e. D_2 is an $(n-1,k-1,l)$ disjointness matrix so, by the induction hypothesis, $\mathrm{rk}_2(D_2)\ge\binom{n-1}{\min(k-1,l)}$. The submatrix of D whose rows include e and whose columns include e is a zero matrix, because all these sets intersect. Thus, the situation is as follows: $D=\begin{pmatrix}0 & D_1\\ D_2 & E\end{pmatrix}$. Because $k\ne l$, it follows that

$$\mathrm{rk}_2(D_n(k,l))\ge \mathrm{rk}_2(D_1)+\mathrm{rk}_2(D_2)\ge\binom{n-1}{\min(k,l-1)}+\binom{n-1}{\min(k-1,l)}=\binom{n}{m},$$

for $m=\min(k,l)$. □

The term $k+l\le n$ is required so that when n gets smaller in the course of the induction, it will still satisfy $k,l\le n$. Although we only need this lower bound for the proof of Theorem 4.3, for completeness we show the upper bound as well.

For the sake of proving the upper bound, it is useful to consider the *inclusion matrix* $I_n(t,l)$, whose rows are labeled by t-sets and whose columns are labeled by l-sets. $I(A,B)=1$ iff $B\subseteq A$. Similarly, denote by $I_n^*(t,l)$ the inclusion matrix whose columns correspond to subsets of size at most l. It is easy to see that $D_n(t,l)=I_n^*(n-t,l)$. Denote also the column space of a matrix M by $\mathrm{Cols}(M)$.

Lemma A.3 (4.1). *Let k, l, n be natural numbers such that $k+l\le n$, and let $m=\min(k,l)$. Then* $\mathrm{rk}_2 D_n(k,l)=\binom{n}{m}$.

Proof. The upper bound for $l\ge k$ is trivial, as the number of rows in $D_n(k,l)$ is $\binom{n}{k}$. Here is a proof for the case $l\le k$ (see Babai and Frankl [2]): The number of l-subsets that contain a specific j-subset and are contained in a specific t-subset is $\binom{t-j}{l-j}$ if the j-tuple is contained in the t-tuple and 0 otherwise, so for every $0\le j\le l\le t$, $I_n(t,l)I_n(l,j)=\binom{t-j}{l-j}I_n(t,j)$, and thereby $\mathrm{Cols}(I_n(t,j))\subseteq\mathrm{Cols}(I_n(t,l))$. Thus, $\mathrm{rk}_{\mathbb{R}}(I_n^*(t,l))=\mathrm{rk}_{\mathbb{R}}(I_n(t,l))\le\binom{n}{l}$. Now, because $l\le n-k$, we can take $t=n-k$, and get:

$$\mathrm{rk}_2(D_n(k,l))\le\mathrm{rk}_{\mathbb{R}}(D_n(k,l))=\mathrm{rk}_{\mathbb{R}}(I_n^*(n-k,l))\le\binom{n}{l}.\quad\square$$

Remark. All arguments in the appendix hold for any field $\mathbb{F}_p$ where p is prime, and not only $\mathbb{F}_2$.

Acknowledgment. The authors are grateful to Noga Alon for enlightening comments, fruitful discussions, and exposing them to relevant literature.

References

[1] Alon, N., L. Babai, H. Suzuki. 1991. Multilinear polynomials and Frankl-Ray-Chaudhuri-Wilson type intersection theorems. *J. Combin. Theory A* **58** 165–180.
[2] Babai, L., P. Frankl. 1992. Linear algebra methods in combinatorics. Preliminary Version 2. Department of Computer Science, The University of Chicago, Chicago, IL.
[3] Bartholdi, J. J. 1982. A good matrix is hard to find. *Oper. Res. Lett.* **5** 190–193.
[4] Cheng, C. K., T. C. Hu. 1990. Ancestor tree for arbitrary multi-terminal cut functions. R. Kannan, W. R. Pulleyblank, eds. *Proc. 1st Integer Programming Combin. Optim. Conf.* University of Waterloo Press, Waterloo, Ontario, Canada, 115–127.
[5] Eilam-Tzoreff, T. 1996. Disjoint paths in graphs with additional constraints. Ph.D. dissertation, Tel Aviv University, Tel Aviv, Israel.
[6] Frankl, P. 1982. An extremal problem for two families of sets. *Eur. J. Combin.* **3** 125–127.
[7] Frankl, P. 1990. Intersection theorems and mod p rank of inclusion matrices. *J. Combin. Theory Series A* **54** 85–94.
[8] Frumkin, A., A. Yakir. 1990. Rank of inclusion matrices and modular representation theory. *Israel J. Math.* **71**(3) 309–320.
[9] Gomory, R. E., T. C. Hu. 1961. Multi-terminal network flows. *J. SIAM* **9** 551–570.
[10] Hartvigsen, D. 2001. Compact representation of cuts. *SIAM J. Discrete Math.* **14**(1) 49–66.
[11] Hassin, R. 1988. Solution bases of multiterminal cut problems. *Math. Oper. Res.* **13** 535–542.
[12] Hassin, R. 1990. An algorithm for computing maximum solution bases. *Oper. Res. Lett.* **9** 315–318.
[13] Hassin, R. 1991. Multiterminal xcut problems. *Ann. Oper. Res.* **33** 215–225.
[14] Jukna, S. 2002. *Extremal Combinatorics*. Springer, Berlin, Germany.
[15] Linial, N., B. L. Rothschild. 1981. Incidence matrices of subsets—A rank formula. *SIAM J. Algebraic Discrete Methods* **2**(4) 333–340.
[16] Lovász, L. 1979. Topological and algebraic methods in graph theory. *Graph Theory and Related Topics*. Academic Press, New-York, London, UK, 1–14.
[17] Parekh, O. 2003. Forestation in hypergraphs: Linear k-trees. *Electronic J. Combin.* **10**(1).
[18] Razborov, A. A. 1987. Lower bounds on the size of bounded-depth network over a complete basis with logical addition. *Math. Notes Acad. Sci. USSR* **41**(4) 333–338.
[19] Wilson, R. M. 1990. A diagonal form for the incidence matrices of t-subsets vs. k-subsets. *Eur. J. Combin.* **11** 609–615.