

The Parameterized Complexity of the Unique Coverage Problem^{*}

Hannes Moser^{1**}, Venkatesh Raman², Somnath Sikdar²

¹ Institut für Informatik, Friedrich-Schiller-Universität Jena,
Ernst-Abbe-Platz 2, D-07743 Jena, Germany
moser@minet.uni-jena.de

² The Institute of Mathematical Sciences,
C.I.T Campus, Taramani, Chennai 600113, India
{vraman,somnath}@imsc.res.in

Abstract. We consider the parameterized complexity of the UNIQUE COVERAGE problem: given a family of sets and a parameter k , we ask whether there exists a subfamily that covers at least k elements exactly once. This NP -complete problem has applications in wireless networks and radio broadcasting and is also a natural generalization of the well-known MAX CUT problem. We show that this problem is fixed-parameter tractable with respect to the parameter k . That is, for every fixed k , there exists a polynomial-time algorithm for it. One way to prove a problem fixed-parameter tractable is to show that it is kernelizable. To this end, we show that if no two sets in the input family intersect in more than c elements there exists a problem kernel of size k^{c+1} . This yields a k^k kernel for the UNIQUE COVERAGE problem, proving fixed-parameter tractability. Subsequently, we show a 4^k kernel for this problem. However a more general weighted version, with costs associated with each set and profits with each element, turns out to be much harder. The question here is whether there exists a subfamily with total cost at most a prespecified budget B such that the total profit of uniquely covered elements is at least k , where B and k are part of the input. In the most general setting, assuming real costs and profits, the problem is not fixed-parameter tractable unless $P = NP$. Assuming integer costs and profits we show the problem to be $W[1]$ -hard with respect to B as parameter (that is, it is unlikely to be fixed-parameter tractable). However, under some reasonable restriction, the problem becomes fixed-parameter tractable with respect to both B and k as parameters.

1 Introduction

In this paper, we consider the parameterized complexity of the UNIQUE COVERAGE problem. This problem was introduced by Demaine et al. [2] as a natural maximization version of SET COVER and has applications in several areas including wireless networks and radio broadcasting. UNIQUE COVERAGE is

^{*} Supported by a DAAD-DST exchange program, D/05/57666.

^{**} Supported by the Deutsche Forschungsgemeinschaft, project ITKO (iterative compression for solving hard network problems), NI 369/5.

defined as follows. Given a ground set $\mathcal{U} = \{1, 2, \dots, n\}$, a family of subsets $\mathcal{S} = \{S_1, \dots, S_t\}$ of \mathcal{U} and a positive integer k , we ask whether there exists a subcollection $\mathcal{S}' \subseteq \mathcal{S}$ such that at least k elements are covered uniquely by the members in \mathcal{S}' . An element is covered uniquely if it appears in exactly one set of \mathcal{S}' . The optimization version requires to maximize the number of uniquely covered elements.

The weighted version of UNIQUE COVERAGE is called BUDGETED UNIQUE COVERAGE and is defined as follows. Given a ground set $\mathcal{U} = \{1, 2, \dots, n\}$, a profit p_i for each element $i \in \mathcal{U}$, a family of subsets \mathcal{S} of \mathcal{U} , a cost c_i for each set $S_i \in \mathcal{S}$, a budget B and a positive integer k , we ask whether there exists a subset $\mathcal{S}' \subseteq \mathcal{S}$ such that the total cost of \mathcal{S}' is at most B and such that the profit of the uniquely covered elements is at least k . The optimization version asks for a subset \mathcal{S}' of total cost at most B such that the profit of uniquely covered elements is maximized.

The original motivation for this problem is a real-world application arising in wireless networks [2]. Assume that we are given a map of the densities of mobile clients along with a set of possible base stations, each with a specified building cost and a specified coverage region. The goal is to choose a set of base stations, subject to a budget on the total building cost, in order to maximize the density of served clients. The difficult aspect of this problem is the interference between base stations. A mobile client's reception is better when it is within the range of a few base stations. An ideal situation is when every mobile client is within the range of exactly one base station. This is the situation modelled by the BUDGETED UNIQUE COVERAGE problem. The UNIQUE COVERAGE problem is closely related to a single "round" of the RADIO BROADCAST problem [1]. For more about this relation, see Demaine et al.'s work [2].

One can also view the UNIQUE COVERAGE problem as a generalization of the MAX CUT problem [2]. The input to the MAX CUT problem consists of a graph $G = (V, E)$ and the goal is to find a cut (T, T') , where $\emptyset \neq T \subset V$ and $T' = V \setminus T$, that maximizes the number of edges with one endpoint in T and the other endpoint in T' . Let \mathcal{U} denote the set of edges of G and for each vertex $v \in V$ define $S_v = \{e \in E : e \text{ is incident to } v\}$. Finally let $\mathcal{S} = \cup_{v \in V} \{S_v\}$. Then G has a cut (T, T') with at least k edges across it if and only if $\mathcal{S}' = \cup_{v \in T} \{S_v\}$ uniquely covers at least k elements of the ground set.

Known Results. (BUDGETED) UNIQUE COVERAGE was introduced by Demaine et al. [2]. They have considered the approximability of this problem. On the positive side, they give an $\Omega(1/\log n)$ -approximation for BUDGETED UNIQUE COVERAGE. Moreover, if the ratio between the maximum cost of a set and the minimum profit of an element is bounded by B , then there exists an $\Omega(1/\log B)$ -approximation. Concerning approximation hardness, they show that UNIQUE COVERAGE is hard to approximate to within a factor of $O(1/\log^c n)$ for some constant $0 < c \leq 1$, and they strengthen this inapproximability to $O(1/\log n)$ based on a hardness hypothesis for BALANCED BIPARTITE INDEPENDENT SET.

UNIQUE COVERAGE (<i>Parameter: k</i>)	<i>Result</i>	Sect.
Each element occurs in at most b sets	$(k - 1)b$ kernel	3.1
Intersection size bounded by c	k^{c+1} kernel	3.2
General	4^k kernel	3.3
Each set of size at most b	2^{b+k} kernel	3.3
<hr/>		
BUDGETED UNIQUE COVERAGE		
Arbitrary costs/profits (pars. B and k)	Not FPT (unless $P = NP$)	4.1
Integer weights (par. B)	$W[1]$ -hard	4.1
Integer weights (intersection number $f(k)$; pars. B and k)	$O^*((B \cdot 2^{f(k)})^{B+k})$ -time algo.	4.2
Integer weights (pars. B and k)	Open	

Fig. 1. Main results in this paper.

Our Results. In this paper, we give first-time results on the parameterized complexity of UNIQUE COVERAGE and BUDGETED UNIQUE COVERAGE. Compared to the related SET COVER problem, which is $W[2]$ -complete with respect to the number of sets in the solution as parameter³, UNIQUE COVERAGE becomes fixed-parameter tractable with respect to the number of uniquely covered elements. In other words, the number of uniquely covered elements seems to be a good parameter in order to exploit and reveal the inherent structure of coverage problems in general. Our results indicate that the budgeted variant, BUDGETED UNIQUE COVERAGE, is a much harder problem. More specifically, we show the following.

We show that a special case of UNIQUE COVERAGE where any two sets in the input family intersect in at most c elements is fixed-parameter tractable by demonstrating a polynomial kernel of size k^{c+1} . This leads to a problem kernel of size k^k in the general case, proving that UNIQUE COVERAGE is fixed-parameter tractable. However, the general case can be improved using results from extremal combinatorics on strong systems of distinct representatives to obtain a 4^k kernel. For the BUDGETED UNIQUE COVERAGE problem there are several variants. If the profits and costs are allowed to be arbitrary positive real numbers, then BUDGETED UNIQUE COVERAGE, with parameters B and k , is not fixed-parameter tractable unless $P = NP$. If we restrict the costs and profits to be positive integers and parameterize by B , then the problem is $W[1]$ -hard. In the case when the number of sets intersecting any given set of the input family is bounded by a function of k , the problem is fixed-parameter tractable with parameters B and k . The main results of this paper along with the relevant sections in which they appear are depicted in Figure 1.

³ This can be shown by a reduction from DOMINATING SET [3, 6].

2 Preliminaries

We briefly introduce the necessary concepts concerning parameterized complexity. A parameterized problem is a subset of $\Sigma^* \times \mathbb{N}$, where Σ is a finite alphabet and \mathbb{N} is the set of natural numbers. An instance of a parameterized problem is therefore a pair (I, k) , where k is the parameter. In the framework of parameterized complexity, the running time of an algorithm is viewed as a function of two quantities: the size of the problem instance *and* the parameter. A parameterized problem is said to be *fixed parameter tractable* (FPT) if there exists an algorithm for the problem with running time $f(k) \cdot |I|^{O(1)}$, where f is a computable function only depending on k .

A common method to prove that a problem is fixed-parameter tractable is to provide data reduction rules that lead to a problem kernel. A *data reduction rule* is a polynomial-time algorithm which takes a problem instance (I, k) and either

- outputs YES or NO according as (I, k) is a YES or a NO-instance, or
- replaces (I, k) by an equivalent instance (I', k') such that $|I'| \leq |I|$ and $k' \leq k$,

where two problem instances (I, k) and (I', k') are *equivalent* if they are both YES-instances or both NO-instances. An instance to which none of a given set of data reduction rules applies is called *reduced* with respect to this set of rules. A parameterized problem is said to have a *problem kernel* if the resulting reduced instance has size $f(k)$ for a function f depending only on k . If a parameterized problem has a kernel, then it is clearly fixed-parameter tractable. Simply use brute-force on the kernel to decide whether the given instance is a YES-instance or not.

A parameterized problem π_1 is *fixed-parameter reducible* to a parameterized problem π_2 if there exist functions $f, g : \mathbb{N} \rightarrow \mathbb{N}$, $\Phi : \Sigma^* \times \mathbb{N} \rightarrow \Sigma^*$ and a polynomial $p(\cdot)$ such that for any instance (I, k) of π_1 , $(\Phi(I, k), g(k))$ is an instance of π_2 computable in time $f(k) \cdot p(|I|)$ and $(I, k) \in \pi_1$ if and only if $(\Phi(I, k), g(k)) \in \pi_2$. The basic complexity class for fixed-parameter intractability is $W[1]$ as there is strong evidence that $W[1]$ -hard problems are not fixed-parameter tractable [3]. To show that a problem is $W[1]$ -hard, one needs to exhibit a fixed-parameter reduction from a known $W[1]$ -hard problem to the problem at hand. For more on parameterized complexity see [3, 5].

We write $O^*(f(k))$ to denote a running time of $O(f(k) \cdot \text{poly}(n, k))$, where n is the input size and k is the parameter. That is, we use the $O^*(\cdot)$ notation to suppress polynomial factors in the running time.

3 The Unique Coverage Problem

Let $(\mathcal{U} = \{1, 2, \dots, n\}, \mathcal{S} = \{S_1, S_2, \dots, S_m\}, k)$ be an instance of UNIQUE COVERAGE. Apply the following data reduction rules on $(\mathcal{U}, \mathcal{S}, k)$ until no longer applicable.

R1 If there exists $S_i \in \mathcal{S}$ such that $|S_i| \geq k$, then the given instance is a YES-instance.

R2 If there exists $S_1, S_2 \in \mathcal{S}$ such that $S_1 = S_2$, then delete S_1 .

These reduction rules are obviously correct. In the following we always assume that the given instance of UNIQUE COVERAGE is reduced with respect to the above rules.

3.1 Bounded Number of Occurrences

We begin with the simple case where each element $e \in \mathcal{U}$ is contained in at most b sets of \mathcal{S} . A special case of this situation is MAX CUT ($b = 2$).

Lemma 1. *If each element $e \in \mathcal{U}$ occurs in at most b sets of \mathcal{S} then the UNIQUE COVERAGE problem admits a kernel of size $b(k - 1)$.*

Proof. Find a maximal collection \mathcal{S}' of pairwise disjoint sets in \mathcal{S} . If $|\cup_{S \in \mathcal{S}'} S| \geq k$, we are done. Therefore assume $|\cup_{S \in \mathcal{S}'} S| \leq k - 1$. Since every set in $\mathcal{S} - \mathcal{S}'$ intersects some set in \mathcal{S}' and since every element occurs in at most b sets in \mathcal{S} , we have $|\mathcal{S} - \mathcal{S}'| \leq (k - 1)(b - 1)$. But $|\mathcal{S}'| \leq k - 1$ and so $|\mathcal{S}| \leq b(k - 1)$. \square

The proof of Lemma 1 applies a proof principle which is a basis for the proof of the following more complicated case.

3.2 Bounded Intersection Size

Consider the situation where for all $S_i, S_j \in \mathcal{S}$ we have $|S_i \cap S_j| \leq c$, for some constant c . In this case we say that the problem instance has *bounded intersection size* c and show that the problem admits a polynomial kernel of size $O(k^{c+1})$. First consider the case when $|S_i \cap S_j| \leq 1$.

Lemma 2. *Suppose that for all $S_i, S_j \in \mathcal{S}$, $i \neq j$, we have $|S_i \cap S_j| \leq 1$. If an element $e \in \mathcal{U}$ is covered by at least $k + 1$ sets, then one can obtain a solution covering k elements uniquely in polynomial time.*

Proof. Suppose an element $e \in \mathcal{U}$ is covered by the sets S_1, \dots, S_{k+1} . Then by reduction rule R2, at most one of these sets can have size 1. The remaining k sets uniquely cover at least one element each. \square

One can now easily obtain a kernel of size k^2 for the case when the intersection size is at most 1.

Lemma 3. *Suppose that for all $S_i, S_j \in \mathcal{S}$, $|S_i \cap S_j| \leq 1$. If $|\mathcal{S}| \geq k^2$, then there exists $\mathcal{T} \subseteq \mathcal{S}$ that covers at least k elements uniquely.*

Proof. Greedily find a maximal collection $\mathcal{S}' = \{S_1, \dots, S_p\}$ of pairwise disjoint sets in \mathcal{S} . Note that if $|\cup_{S_i \in \mathcal{S}'} S_i| \geq k$, then we are done. Therefore assume, $|\cup_{S_i \in \mathcal{S}'} S_i| \leq k - 1$ (this also implies $p \leq k - 1$). Since $|\mathcal{S}| \geq k^2$, and since every set in \mathcal{S} intersects with at least one set in \mathcal{S}' , by the pigeonhole principle there exists an

element $e \in \cup_{S \in \mathcal{S}'} S$ such that at least $k+1$ sets T_1, \dots, T_{k+1} in $\mathcal{S} - \{S_1, \dots, S_p\}$ contain e . For otherwise, each element in $\cup_{S \in \mathcal{S}'} S$ is contained in at most k sets of $\mathcal{S} \setminus \mathcal{S}'$, which implies that $|\mathcal{S}| \leq (k-1)k + p < k^2$, a contradiction. By Lemma 2, this collection $\mathcal{T} = \{T_1, \dots, T_{k+1}\}$ of $k+1$ sets uniquely covers at least k elements. \square

Next, we generalize these observations to the case when $|S_i \cap S_j| \leq c$, for some constant c .

Theorem 1. *Suppose that for all $S_i, S_j \in \mathcal{S}$ we have $|S_i \cap S_j| \leq c$, for some positive constant c . If $|\mathcal{S}| \geq k^{c+1}$ then there exists $\mathcal{T} \subseteq \mathcal{S}$ that covers k elements uniquely.*

Proof. By induction on c . For $c = 1$, this follows from Lemma 3. Assume the theorem to hold for $c > 1$. Greedily obtain a maximal collection $\mathcal{S}' = \{S_1, \dots, S_p\}$ of pairwise disjoint sets. If $|\cup_{S_i \in \mathcal{S}'} S_i| \geq k$ then we are done. Therefore assume $|\cup_{S \in \mathcal{S}'} S| \leq k-1$ (this also implies $p \leq k-1$). Since $|\mathcal{S}| \geq k^{c+1}$, and since every set in \mathcal{S} intersects with at least one set in \mathcal{S}' , there exists $e \in \cup_{S \in \mathcal{S}'} S$ such that at least $k^c + 1$ sets in $\mathcal{S} - \{S_1, \dots, S_p\}$ contain e . For otherwise, $|\mathcal{S}| \leq (k-1)k^c + p < k^{c+1}$, a contradiction. Let T_1, \dots, T_{k^c+1} be some $k^c + 1$ such sets. Delete e from each of these sets. We obtain at least k^c nonempty distinct sets T'_1, \dots, T'_{k^c} (there is at most one set consisting of the element e only which is deleted in this process). Note that any two of these sets intersect in at most $c-1$ elements. By induction hypothesis, there exists a collection $\mathcal{T}' \subseteq \{T'_1, \dots, T'_{k^c}\}$ that uniquely covers at least k elements, and thus there exists a collection $\mathcal{T} \subseteq \{T_1, \dots, T_{k^c}\}$ that uniquely covers at least k elements (just take the solution for \mathcal{T}' and add e to every set in it). This proves the theorem. \square

Corollary 1. *UNIQUE COVERAGE admits a kernel of size k^{c+1} for bounded intersection size c .*

Note that $c \leq k-1$ and therefore for the general case we have a kernel of size k^k .

Corollary 2. *The UNIQUE COVERAGE problem is fixed-parameter tractable and admits a problem kernel of size k^k .*

An algorithm that checks all possible subsets of a family of size k^k to see whether any of them uniquely covers at least k elements is an FPT algorithm with time complexity $O^*(2^{(k^k)})$. But note that we can assume without loss of generality that every set in the solution covers at least one element uniquely. Thus it suffices to check whether subfamilies of size at most k uniquely cover at least k elements. This can be done in time $O^*(k^{k^2}) = O^*(2^{k^2 \log k})$. However, this kernelization result is tailored especially for the bounded intersection size case. It turns out that a much better kernel can be obtained for the general case.

3.3 General Case

We now show that UNIQUE COVERAGE has a kernel of size 4^k using a result on strong systems of distinct representatives. Given a family of sets $\mathcal{S} = \{S_1, \dots, S_m\}$, a *system of distinct representatives* for \mathcal{S} is an m -tuple (x_1, \dots, x_m) where the elements x_i are distinct and $x_i \in S_i$ for all $i = 1, 2, \dots, m$. Such a system is *strong* if we additionally have $x_i \notin S_j$ for all $i \neq j$. The next theorem due to Füredi and Tuza appears in Jukna's textbook [4].

Theorem 2. *In any family of more than $\binom{r+s}{s}$ sets of cardinality at most r , at least $s + 2$ of its members have a strong system of distinct representatives.*

Given an instance $(\mathcal{U} = \{1, \dots, n\}, \mathcal{S} = \{S_1, \dots, S_m\}, k)$ of UNIQUE COVERAGE, put $r = k - 1$ and $s = k$ in the statement of the above theorem and we have a kernel of size $\binom{2k-1}{k-1} \leq \binom{2k}{k} \leq 2^{2k}$.

Corollary 3. UNIQUE COVERAGE admits a problem kernel of size 4^k .

Note that this implies that there is an $O^*(4^{k^2})$ time FPT algorithm for the UNIQUE COVERAGE problem.

For the case where each set of the input family has size at most b , for some constant b , there is a better kernel. By Theorem 2, if there exists at least $\binom{b+k}{k}$ sets in the input family, then there exists at least k sets with a strong system of distinct representatives.

Corollary 4. *If each set $S \in \mathcal{S}$ has size at most b then the UNIQUE COVERAGE problem has a kernel of size $O(2^{b+k})$.*

4 The Budgeted Unique Coverage Problem

In this section we consider the BUDGETED UNIQUE COVERAGE problem where each set in the input family has a cost and each element in the universe has a profit; the goal is to decide whether there exists a subcollection of total cost at most B that uniquely covers elements of total profit at least k . By parameterizing on k or B or both we obtain different parameterized versions of this decision question.

4.1 Hardness Results

We first consider the BUDGETED MAX CUT problem which is a specialization of the BUDGETED UNIQUE COVERAGE problem. An instance of this problem is an undirected graph $G = (V, E)$ with a cost function $c : V \rightarrow \mathbb{R}^+$ on the vertex set and a profit function $p : E \rightarrow \mathbb{R}^+$ on the edge set; positive real numbers B and k . The question is whether there exists a cut (T, T') such that the total cost of the vertices in T is at most B and the total profit of the edges crossing the cut is at least k .

We first show that the BUDGETED MAX CUT problem with arbitrary positive real costs and profits is probably not FPT.

Lemma 4. *The BUDGETED MAX CUT problem with arbitrary positive costs and profits with parameters B and k is not FPT, unless $P = NP$.*

Proof. Suppose there exists an algorithm for the BUDGETED MAX CUT problem (with arbitrary positive costs and profits) with running time $O(f(k, B) \cdot \text{poly}(n))$. We will use this to solve the decision version of MAX CUT in polynomial time. Let $(G = (V, E), k)$ be an instance of the MAX CUT problem, where $|V| = n$. Assign each vertex of the input graph cost $1/n$ and each edge profit $1/k$. Let the budget $B = 1/2$ and the profit $k' = 1$. Clearly, G has a maximum cut of size at least k iff there exists $S \subseteq V$ of total cost at most B such that the total profit of the edges crossing the cut $(S, V - S)$ is at least k' . And this can be answered in time $O(f(1, 1/2) \cdot p(|V|))$, implying $P = NP$. \square

Theorem 3. *The BUDGETED UNIQUE COVERAGE problem with arbitrary positive costs and profits is not FPT, unless $P = NP$.*

Henceforth by the ‘budgeted’ version we mean the case when the costs and profits are positive integers. We next show that the BUDGETED MAX CUT problem parameterized by the budget B alone is $W[1]$ -hard.

Lemma 5. *The BUDGETED MAX CUT problem parameterized by the budget is $W[1]$ -hard.*

Proof. To show $W[1]$ -hardness, we exhibit a fixed-parameter reduction from the INDEPENDENT SET problem to the BUDGETED MAX CUT problem with unit costs and profits. Let $(G = (V, E), B)$ be an instance of INDEPENDENT SET with $|V| = n$. For every vertex $u \in V$ add $|V| - 1 - \text{deg}(u)$ new vertices and connect them to u . Call the resulting graph G' . Note that every vertex $u \in G$ has degree $|V| - 1$ in G' . We let $(G' = (V', E'), B, k = B(n - 1))$ be the instance of BUDGETED MAX CUT.

Claim. G has an independent set of size B iff G' has a cut $(S, V' - S)$ such that $|S| = B$ and at least $k = B(n - 1)$ edges lie across it.

If G has an independent set S of size B , then clearly S is independent in G' . The cut $(S, V' - S)$ does indeed have $B(n - 1)$ edges crossing it, as every vertex of S has degree $n - 1$. Next suppose that G' has a cut $(S, V' - S)$ such that $|S| = B$ and $B(n - 1)$ edges cross the cut. Note that every vertex in S must be a vertex from G . Otherwise the cut cannot have $B(n - 1)$ edges crossing it. Suppose two vertices u and v in S are adjacent. Then both u and v contribute less than $n - 1$ edges to the cut. Since each vertex in S contributes at most $n - 1$ edges to the cut, the number of edges crossing the cut must be less than $B(n - 1)$, a contradiction. Hence S is independent in G' and hence G has an independent set of size B . \square

Since the BUDGETED UNIQUE COVERAGE problem is a generalization of BUDGETED MAX CUT we have

Theorem 4. *The BUDGETED UNIQUE COVERAGE problem parameterized by the budget B is $W[1]$ -hard.*

4.2 A Fixed-Parameter Tractability Result

In this subsection, we give an FPT algorithm for BUDGETED UNIQUE COVERAGE, when B and k are parameters, assuming that for every set S in the input family the number of sets with a non-empty intersection with S is at most some function of k . This is a natural situation in real-world applications; for example, in wireless networks. For the BUDGETED MAX CUT problem, for instance, every set is intersected by at most $k - 1$ sets.

Let $(\mathcal{U} = \{1, \dots, n\}, \mathcal{S} = \{S_1, \dots, S_m\}, c, p, B, k)$ be an instance of the BUDGETED UNIQUE COVERAGE problem where $c : \mathcal{S} \rightarrow \mathbb{N}$ and $p : \mathcal{U} \rightarrow \mathbb{N}$. For $\mathcal{T} \subseteq \mathcal{S}$, define $c(\mathcal{T}) = \sum_{S \in \mathcal{T}} c(S)$ and $p(\mathcal{T})$ to be the total profit of the elements uniquely covered by \mathcal{T} . If $S_i \in \mathcal{S}$, define $N[S_i]$ to be the set of all members of \mathcal{S} that have a nonempty intersection with S_i . We can without loss of generality assume that $c(S_i) \leq B$ and $|S_i| \leq k - 1$ for all $1 \leq i \leq m$. In what follows, we assume that for all $S_i \in \mathcal{S}$, we have $|N[S_i]| \leq f(k)$ for some function f .

The FPT algorithm that we describe here builds the solution in stages. Note that if we decide to include a set S in the solution, there is no way of deciding how many elements S covers uniquely unless we make choices for each set in $N[S]$. To get around this, the algorithm, at any stage, decides whether or not to include a subfamily $\mathcal{A} \subseteq N[S]$ for some set S . If it includes \mathcal{A} in the solution, then it automatically *excludes* $N[S] \setminus \mathcal{A}$ from it. The current solution is a pair $(\mathcal{T}, \mathcal{T}')$, where $\mathcal{T}, \mathcal{T}' \subseteq \mathcal{S}$ and $\mathcal{T} \cap \mathcal{T}' = \emptyset$. The sets included by the algorithm in the solution till the current stage are in \mathcal{T} ; those excluded from the solution are in \mathcal{T}' .

Call a pair $(\mathcal{T}, \mathcal{T}')$ a feasible solution for an instance of BUDGETED UNIQUE COVERAGE if $\mathcal{T}' = \mathcal{S} - \mathcal{T}$, $c(\mathcal{T}) \leq B$ and $p(\mathcal{T}) \geq k$. A pair $(\mathcal{T}, \mathcal{T}')$ is a *partial solution* if $\mathcal{T}, \mathcal{T}' \subseteq \mathcal{S}$ and $\mathcal{T} \cap \mathcal{T}' = \emptyset$. A partial solution $(\mathcal{T}, \mathcal{T}')$ can be *extended* to a feasible solution if there exist $\mathcal{X}, \mathcal{X}' \subseteq \mathcal{S} - (\mathcal{T} \cup \mathcal{T}')$ such that $\mathcal{X} \cap \mathcal{X}' = \emptyset$ and $(\mathcal{T} \cup \mathcal{X}, \mathcal{T}' \cup \mathcal{X}')$ is a feasible solution. A partial solution $(\mathcal{T}, \mathcal{T}')$ is *strong* if for each set $S_i \in \mathcal{T}$, $N[S_i] \subseteq \mathcal{T} \cup \mathcal{T}'$. Given a strong partial solution $(\mathcal{T}, \mathcal{T}')$, let $\mathcal{U}_1, \dots, \mathcal{U}_t$ be a partition of $\mathcal{S} - (\mathcal{T} \cup \mathcal{T}')$ according to costs. That is, all members in any set \mathcal{U}_i have the same cost c_i and for all $i \neq j$, $c_i \neq c_j$. Note that $t \leq B$. For each \mathcal{U}_i , let U_i^{max} denote a member of \mathcal{U}_i with maximum total profit.

Lemma 6. *Let $(\mathcal{T}, \mathcal{T}')$ be a strong partial solution and let $\mathcal{U}_1, \mathcal{U}_2, \dots, \mathcal{U}_t$ be a partition of $\mathcal{S} - (\mathcal{T} \cup \mathcal{T}')$ according to costs. Suppose $(\mathcal{T}, \mathcal{T}')$ can be extended to a feasible solution by adding a member of \mathcal{U}_i to \mathcal{T} . Then there exists an extension of $(\mathcal{T}, \mathcal{T}')$ into a feasible solution such that $\mathcal{T} \cap N[U_i^{max}] \neq \emptyset$.*

Proof. Suppose $(\mathcal{T}, \mathcal{T}')$ can be extended to a feasible solution $(\mathcal{X}, \mathcal{X}')$ by adding a member $U \in \mathcal{U}_i$ to \mathcal{T} and that $\mathcal{X} \cap N[U_i^{max}] = \emptyset$. This means $N[U_i^{max}] \subseteq \mathcal{X}'$. Remove U from \mathcal{X} and replace it by U_i^{max} . Note that every element of U_i^{max} is uniquely covered and that the total profit of these newly uniquely covered elements is at least as that of those covered by U . Since $c(U) = c(U_i^{max})$, the new solution continues to be feasible. \square

One can use Lemma 6 to develop an FPT algorithm with time complexity $O^*((B \cdot 2^{f(k)})^{B+k})$. Suppose there exists a feasible solution to the given input instance. We start with a strong feasible solution $(\mathcal{T} = \emptyset, \mathcal{T}' = \emptyset)$. Partition the input family \mathcal{S} according to costs into the subfamilies $\mathcal{U}_1, \dots, \mathcal{U}_t$. Note that $t \leq B$. Since there exists a feasible solution, it has to include a set from one of the subfamilies \mathcal{U}_i . For each choice of a subfamily, Lemma 6 assures us that it is sufficient to consider a set S in the subfamily which maximizes profit. We consider all possible bipartitions $(\mathcal{A}, \mathcal{A}')$ of $N[S]$ such that $\mathcal{A} \neq \emptyset$ and each member in \mathcal{A} uniquely covers at least one element. For each such bipartition $(\mathcal{A}, \mathcal{A}')$, set $\mathcal{T} \leftarrow \mathcal{T} \cup \mathcal{A}$ and $\mathcal{T}' \leftarrow \mathcal{T}' \cup \mathcal{A}'$. Since by our assumption, $|N[S]| \leq f(k)$, there are at most $2^{f(k)}$ such bipartitions. This gives an initial branching factor of $B \cdot 2^{f(k)}$.

We then recurse using Lemma 6. In order to recurse, we must first ensure that the current partial solution is strong. We achieve this by considering all possible bipartitions of $N[T] - (\mathcal{T} \cup \mathcal{T}')$ for all sets $T \in \mathcal{T}$ for which $N[T] - (\mathcal{T} \cup \mathcal{T}') \neq \emptyset$. As before, we are interested in bipartitions $(\mathcal{A}, \mathcal{A}')$ which have the property that each set in $\mathcal{T} \cup \mathcal{A}$ uniquely covers at least one element. For each such bipartition $(\mathcal{A}, \mathcal{A}')$, we set $\mathcal{T} \leftarrow \mathcal{T} \cup \mathcal{A}$ and $\mathcal{T}' \leftarrow \mathcal{T}' \cup \mathcal{A}'$. There are at most 2^k such bipartitions and for each bipartition, we either increase the cost of the solution or total profit of uniquely covered elements by at least 1. If at any stage of recursion, we find that there is no subfamily \mathcal{U}_i such that for $U \in \mathcal{U}_i$, $c(U) \leq B - c(\mathcal{T})$, we abort that branch. If $p(\mathcal{T}) \geq k$, at any stage, we halt and output YES. The overall depth of the recursion tree is at most $B + k$ and the branching factor is at most $B \cdot 2^{f(k)}$. The overall time complexity is therefore $O^*((B \cdot 2^{f(k)})^{B+k})$. If the algorithm does not return a solution then we can safely conclude that the given instance is a NO-instance.

Theorem 5. *Suppose $(\mathcal{U}, \mathcal{S}, c, p, B, k)$ is an instance of the BUDGETED UNIQUE COVERAGE problem where for every set $S \in \mathcal{S}$, we have $|N[S]| \leq f(k)$. Then there is an algorithm with time complexity $O^*((B \cdot 2^{f(k)})^{B+k})$ for this problem.*

The BUDGETED MAX CUT problem is a special case where $|N[S]| \leq k - 1$ for all $S \in \mathcal{S}$, and the following corollary is immediate.

Corollary 5. *The BUDGETED MAX CUT problem with positive integer costs and profits is fixed-parameter tractable when parameterized by B and k . There is an algorithm with time complexity $O^*((B \cdot 2^k)^{B+k})$ for this problem.*

5 Concluding Remarks

In this paper, we considered the parameterized complexity of the UNIQUE COVERAGE problem. There are several directions in which to proceed. Firstly, the reduction rules that we give are almost trivial and the kernel that we obtain is exponential in k . Kernelization is a very important topic in the design of FPT algorithms and the challenge is to devise reduction rules to obtain a polynomial (linear?) kernel or prove that no such kernel exists under some plausible

complexity-theoretic assumption. Are there reduction rules that lead to a better problem kernel? In particular, is there a polynomial kernel for the UNIQUE COVERAGE problem?

At this point, all we can show is that with respect to a broader set of reduction rules, which we do not state here, the kernel size is at least $\Omega(2^k/\sqrt{k/2})$. The following example illustrates this situation. Let $\mathcal{U} = \{1, 2, \dots, k\}$, $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2$, where \mathcal{S}_1 consists of all subsets of \mathcal{U} of size exactly $\lceil k/2 \rceil + 1$ and \mathcal{S}_2 is some collection of subsets of \mathcal{U} of size at most $k/4$. Note that $|\mathcal{S}_1| = \binom{k}{\lceil k/2 \rceil + 1}$, which by Stirling's approximation is, $\Omega(2^k/\sqrt{k/2})$. If $\mathcal{S}_2 = \emptyset$ then one can show that the given instance is a NO-instance. But we can always produce an $\mathcal{S}_2 \neq \emptyset$ such that the given instance is a YES-instance and such that our reduction rules do not change the size of the input instance. For instance, if we take $\mathcal{S}_2 = \{\{\lceil k/2 \rceil + 2\}, \dots, \{k\}\}$, then this is a YES-instance and we can show that our reduction rules do not alter the size of the input.

Another important question is whether there exists a good branching algorithm for UNIQUE COVERAGE. The algorithm that we gave runs in time $O^*(4^{k^2})$. Finally, is the BUDGETED UNIQUE COVERAGE problem with positive integer costs/profits, with parameters B and k , fixed-parameter tractable?

Acknowledgements We thank Saket Saurabh for pointing out the connection between UNIQUE COVERAGE and strong systems of distinct representatives.

References

1. R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization. *Journal of Computer and System Sciences*, 45:104–126, 1992.
2. Erik D. Demaine, Mohammad T. Hajiaghayi, Uriel Feige, and Mohammad R. Salavatipour. Combination can be hard: approximability of the unique coverage problem. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2006)*, pages 162–171. SIAM, 2006.
3. Rod R. Downey and Michael R. Fellows. *Parameterized Complexity*. Springer Verlag, New York, 1999.
4. Stasys Jukna. *Extremal combinatorics*. Springer-Verlag, Berlin, 2001.
5. Rolf Niedermeier. *An Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006.
6. Azaria Paz and Shlomo Moran. Non deterministic polynomial optimization problems and their approximations. *Theoretical Computer Science*, 15:251–277, 1981.