

Iterative Compression and Exact Algorithms

Fedor V. Fomin* Serge Gaspers† Dieter Kratsch‡
Mathieu Liedloff§ Saket Saurabh¶

Abstract

Iterative Compression has recently led to a number of breakthroughs in parameterized complexity. Here, we show that the technique can also be useful in the design of exact exponential time algorithms to solve NP-hard problems. We exemplify our findings with algorithms for the MAXIMUM INDEPENDENT SET problem, a parameterized and a counting version of d -HITTING SET and the MAXIMUM INDUCED CLUSTER SUBGRAPH problem.

1 Introduction

Iterative Compression is a tool that has recently been used successfully in solving a number of problems in the area of parameterized complexity. This technique was first introduced by Reed et al. to solve the ODD CYCLE TRANSVERSAL problem, where one is interested in finding a set of at most k vertices whose deletion makes the graph bipartite [23]. Iterative compression was used in obtaining faster FPT algorithms for FEEDBACK VERTEX SET, EDGE BIPARTIZATION and CLUSTER VERTEX DELETION on undirected graphs [4, 8, 15, 17]. Recently this technique has led to an FPT algorithm for the DIRECTED FEEDBACK VERTEX SET problem [6], one of the longest open problems in the area of parameterized complexity.

Typically iterative compression algorithms are designed for parameterized minimization problems with a parameter, say k . Such algorithms proceed by iterating the so called *compression step*: given a solution of size

*Department of Informatics, University of Bergen, N-5020 Bergen, Norway. E-mail: fomin@ii.uib.no. Partially supported by the Research Council of Norway.

†Centro de Modelamiento Matemático, Universidad de Chile, 8370459 Santiago de Chile. E-mail: sgaspers@dim.uchile.cl

‡Laboratoire d'Informatique Théorique et Appliquée, Université Paul Verlaine - Metz, 57045 Metz Cedex 01, France. E-mail: kratsch@univ-metz.fr.

§Laboratoire d'Informatique Fondamentale d'Orléans, Université d'Orléans, 45067 Orléans Cedex 2, France. E-mail: liedloff@univ-orleans.fr.

¶Institute of Mathematical Sciences, CIT Campus, Taramani, 600 113 Chennai, India. E-mail: saket@imsc.res.in.

$k + 1$, either compress it to a solution of size k or prove that there is no solution of size k . To obtain a fixed parameter tractable (FPT) algorithm, one has to solve the compression step in time $f(k)n^{O(1)}$, where f is an arbitrary computable function, k is a parameter and n is the length of the input. Technically speaking, almost all the iterative compression based FPT algorithms with parameter k have $f(k) \geq 2^{k+1}$, as they all branch on all partitions (A, D) of a $k + 1$ sized solution S and look for a solution of size k with the restriction that it should contain all elements of A and none of D .

Given the success of iterative compression in designing fixed parameter tractable algorithms, it is natural and tempting to study its applicability in designing exact exponential time algorithms solving computationally hard problems. The goal of the design of such algorithms is to establish algorithms of best possible worst-case running time; at least provably faster than the one of enumerating all prospective solutions, or loosely speaking, *algorithms better than trivial enumeration*. For example, for many NP-hard or #P-hard problems on graphs on n vertices and m edges the prospective solutions are either subsets of vertices (or edges). Thus a trivial algorithm to solve such an NP-hard problem basically enumerates all subsets of vertices (or edges) and finds an optimal solution or all optimal solutions, counts the number of optimal solutions, or enumerates all optimal solutions. This mostly leads to algorithms of time complexity 2^n (or 2^m), modulo some polynomial factors. Our aim is to demonstrate that iterative compression can become a useful tool in the design of exact exponential time algorithms as well; we concentrate on optimization and counting problems.

One simple way to obtain an exact exponential time algorithm from an FPT algorithm is to use the latter for all possible values of the parameter k . In many cases this does not lead to faster exact algorithms. Assuming that the largest (reasonable) value of the parameter k is at least n , using a typical iterative compression based FPT algorithm does not seem to be of much use for constructing an exact exponential time algorithm because we would end up with an algorithm for the compression step having a factor of 2^n (or 2^m) in its worst-case running time; and hence the established algorithm would not be better than trivial enumeration.

There are various well-known techniques to design exact exponential time algorithms for decision, optimization, counting or enumeration problems such as branching algorithms, dynamic programming (over subsets) and inclusion-exclusion algorithms. (We refer the interested reader to [27].) Iterative compression provides combinatorial algorithms based on certain structural properties of the problem (in the compression step). While the improvement in the running time compared to (complicated) branching algorithms is not so impressive, the simplicity and elegance of the arguments allow them to be used in a basic algorithms course. We find this worth emphasising since, despite several exceptions (like the works of Björklund et al.

[1, 2, 19]), the area of exact algorithms is heavily dominated by branching algorithms, in particular, for subset problems. It is very often the case that an (incremental) improvement in the running time of a branching algorithm requires an extensive case analysis, which becomes very technical and tedious. The analysis of such algorithms can also be very complicated and even computer based.

To our knowledge, this paper is the first attempt to use iterative compression outside the domain of FPT algorithms. We exemplify this approach by the following results:

1. We show how to solve MAXIMUM INDEPENDENT SET for a graph on n vertices in time $\mathcal{O}(1.3196^n)$. While the running time of our iterative compression algorithm is higher than the running times of modern branching algorithms [12, 13, 24], this simple algorithm serves as an introductory example to more complicated applications of the method.
2. We obtain algorithms counting the number of minimum hitting sets of a family of sets of an n -element ground set in time $\mathcal{O}(1.7198^n)$, when the size of each set is at most 3 (#MINIMUM 3-HITTING SET). For #MINIMUM 4-HITTING SET we obtain an algorithm of running time $\mathcal{O}(1.8997^n)$. For MINIMUM 4-HITTING SET similar ideas lead to an algorithm of running time $\mathcal{O}(1.8704^n)$. These algorithms are faster than the best algorithms known for these problems so far [11, 22].
3. We provide an algorithm to solve the MAXIMUM INDUCED CLUSTER SUBGRAPH problem in time $\mathcal{O}^*(\phi^n)$ where $\phi = (1 + \sqrt{5})/2 < 1.6181$ is the golden ratio. The only previous algorithm for this problem we are aware of is a very complicated branching algorithm of Wahlström [26] for solving 3-HITTING SET (let us note that MAXIMUM INDUCED CLUSTER SUBGRAPH is a special case of 3-HITTING SET, where every subset is a set of vertices inducing a path of length three), which results in a running time of $\mathcal{O}(1.6278^n)$.

As a byproduct, we provide FPT algorithms for the (k, d) -HITTING SET problem, where the input is a family of sets of size at most d , the parameter is k , and the task is to find a hitting set of size at most k if one exists. For $d = 4$ and $d = 5$ we achieve the respective running times $\mathcal{O}(3.0755^k \cdot n^{\mathcal{O}(1)})$ and $\mathcal{O}(4.0755^k \cdot n^{\mathcal{O}(1)})$, improving the best known algorithms for these cases [11].

2 Maximum Independent Set

MAXIMUM INDEPENDENT SET (MIS) is one of the well studied problems in the area of exact exponential time algorithms and many papers have been written on this problem [12, 13, 24, 25]. It is customary that if we develop a

new method then we first apply it to well known problems in the area. Here, as an introductory example, we consider the NP-complete problem MIS.

MAXIMUM INDEPENDENT SET (MIS): Given a graph $G = (V, E)$ on n vertices, find a maximum independent set of G . An *independent set* of G is a set of vertices $I \subseteq V$ such that no two vertices of I are adjacent in G . A *maximum independent set* is an independent set of maximum size.

It is well-known that I is an independent set of a graph G iff $V \setminus I$ is a vertex cover of G , i.e. every edge of G has at least one end point in $V \setminus I$. Therefore MINIMUM VERTEX COVER (MVC) is the complement of MIS in the sense that I is a maximum independent set of G iff $V \setminus I$ is a minimum vertex cover of G . This fact implies that when designing exponential time algorithms we may equivalently consider MVC. We proceed by defining a compression version of the MVC problem.

COMP-MVC: Given a graph $G = (V, E)$ with a vertex cover $S \subseteq V$, find a vertex cover of G of size at most $|S| - 1$ if one exists.

Note that if we can solve COMP-MVC efficiently then we can solve MVC efficiently by repeatedly applying an algorithm for COMP-MVC as follows. Given a graph $G = (V, E)$ on n vertices with $V = \{v_1, v_2, \dots, v_n\}$, let $G_i = G[\{v_1, v_2, \dots, v_i\}]$ and let C_i be a minimum vertex cover of G_i . By V_i we denote the set $\{v_1, v_2, \dots, v_i\}$. We start with G_1 and put $C_1 = \emptyset$. Suppose that we already have computed C_i for the graph G_i for some $i \geq 1$. We form an instance of COMP-MVC with input graph G_{i+1} and $S = C_i \cup \{v_{i+1}\}$. In this stage we either *compress* the solution S which means that we find a vertex cover S' of G_{i+1} of size $|S| - 1$ and put $C_{i+1} = S'$, or (if there is no such S') we put $C_{i+1} = S$.

Our algorithm is based on the following lemma.

Lemma 1. *Let G_{i+1} and S be given as above. If there exists a vertex cover C_{i+1} of G_{i+1} of size $|S| - 1$, then it can be partitioned into two sets A and B such that*

- (a) $A \subset S$, $|A| \leq |S| - 1$ and A is a minimal vertex cover of $G_{i+1}[S]$.
- (b) $B \subseteq (V_{i+1} \setminus A)$ is a minimum vertex cover of the bipartite graph $G_{i+1}[V_{i+1} \setminus A]$.

Proof. Let C_{i+1} be a vertex cover of G_{i+1} of size $|S| - 1$. Its complement $V_{i+1} \setminus C_{i+1}$ is an independent set. We define $A' = C_{i+1} \cap S$ and $B' = C_{i+1} \setminus A'$. Then A' is a vertex cover of $G_{i+1}[S]$ and $|A'| \leq |S| - 1$. Let $A \subseteq A'$ be a minimal vertex cover of $G_{i+1}[S]$. We define $B = B' \cup (A' \setminus A)$. Since A is a minimal vertex cover of $G_{i+1}[S]$, we have that $S \setminus A$ is an independent set. This in turn implies that $G_{i+1}[V_{i+1} \setminus A]$ is a bipartite graph (with a

bipartition $S \setminus A$ and $V_{i+1} \setminus S$). Finally, since C_{i+1} is a minimum vertex cover of G_{i+1} , we conclude that B is a minimum vertex cover of $G_{i+1}[V_{i+1} \setminus A]$. \square

Lemma 1 implies that the following algorithm solves COMP-MVC correctly.

Step 1: Enumerate all minimal vertex covers of size at most $|S| - 1$ of $G_{i+1}[S]$ as a possible candidate for A .

Step 2: For each minimal vertex cover A find a minimum vertex cover B of the bipartite graph $G_{i+1}[V_{i+1} \setminus A]$ (via the computation of a maximum matching in this bipartite graph [16]).

Step 3: If the algorithm finds a vertex cover $A \cup B$ of size $|S| - 1$ in this way, set $C_{i+1} = A \cup B$, else set $C_{i+1} = S$.

Steps 2 and 3 of the algorithm can be performed in polynomial time, and the running time of Step 1, which is exponential, dominates the running time of the algorithm. To enumerate all maximal independent sets or equivalently all minimal vertex covers of a graph in Step 1, one can use the polynomial-delay algorithm of Johnson et al. [18].

Proposition 2 ([18]). *All maximal independent sets of a graph can be enumerated with polynomial delay.*

For the running time analysis of the algorithm we need the following bounds on the number of maximal independent sets or minimal vertex covers due to Moon and Moser [20] and Byskov [3].

Proposition 3 ([20]). *A graph on n vertices has at most $3^{n/3}$ maximal independent sets.*

Proposition 4 ([3]). *The maximum number of maximal independent sets of size at most k in any graph on n vertices for $k \leq n/3$ is*

$$N[n, k] = \lfloor n/k \rfloor^{(\lfloor n/k \rfloor + 1)k - n} (\lfloor n/k \rfloor + 1)^{n - \lfloor n/k \rfloor k}.$$

Moreover, all such sets can be enumerated in time $\mathcal{O}^*(N[n, k])$.¹

Since

$$\max \left\{ \max_{0 \leq \alpha \leq 3/4} (3^{\alpha n/3}), \max_{3/4 < \alpha \leq 1} (N[\alpha n, (1 - \alpha)n]) \right\} = \mathcal{O}^*(2^{2n/5}),$$

¹Throughout this paper we use a modified big-Oh notation that suppresses all polynomially bounded factors. For functions f and g we write $f(n) = \mathcal{O}^*(g(n))$ if $f(n) = \mathcal{O}(g(n) \text{poly}(n))$, where $\text{poly}(n)$ is a polynomial. Furthermore, since $c^n \cdot \text{poly}(n) = \mathcal{O}((c + \epsilon)^n)$ for any $\epsilon > 0$, we omit polynomial factors in the big-Oh notation every time we round the base of the exponent.

we have that by Propositions 2, 3, and 4, all minimal vertex covers of $G_{i+1}[S]$ of size at most $|S| - 1$ can be listed in time $\mathcal{O}^*(2^{2n/5}) = \mathcal{O}(1.3196^n)$.

Thus, the overall running time of the algorithm solving COMP-MVC is $\mathcal{O}(1.3196^n)$. Since the rounding of the base of the exponent dominates the polynomial factor of the other steps of the iterative compression, we obtain the following theorem.

Theorem 5. *The given algorithm for the problems MAXIMUM INDEPENDENT SET and MINIMUM VERTEX COVER, established by iterative compression, has running time $\mathcal{O}(1.3196^n)$ on graphs of n vertices.*

3 # d -Hitting Set and parameterized d -Hitting Set

The HITTING SET problem is a generalization of VERTEX COVER. Here, given a family of sets over a ground set of n elements, the objective is to hit every set of the family with as few elements of the ground set as possible. We study a version of the hitting set problem where every set in the family has at most d elements.

MINIMUM d -HITTING SET (MHS $_d$): Given a universe V of n elements and a collection \mathcal{C} of subsets of V of size at most d , find a minimum hitting set of \mathcal{C} . A *hitting set* of \mathcal{C} is a subset $V' \subseteq V$ such that every subset of \mathcal{C} contains at least one element of V' .

A counting version of the problem is #MINIMUM d -HITTING SET (#MHS $_d$) that asks for the number of different minimum hitting sets. We denote an instance of #MHS $_d$ by (V, \mathcal{C}) . Furthermore we assume that for every $v \in V$, there exists at least one set in \mathcal{C} containing it.

We show how to obtain an algorithm to solve #MHS $_d$ using iterative compression which uses an algorithm for #MHS $_{d-1}$ as a subroutine. First we define the compression version of the #MHS $_d$ problem.

COMP-# d -HITTING SET: Given a universe V of n elements, a collection \mathcal{C} of subsets of V of size at most d , and a (not necessarily minimum) hitting set $H' \subseteq V$ of \mathcal{C} , find a minimum hitting set \hat{H} of \mathcal{C} and compute the number of all minimum hitting sets of \mathcal{C} .

Lemma 6. *Let $\mathcal{O}^*(a_{d-1}^n)$ be the running time of an algorithm solving #MHS $_{d-1}$, where $a_{d-1} > 1$ is some constant. Then COMP-# d -HITTING SET can be solved in time*

$$\mathcal{O}^* \left(2^{|H'|} a_{d-1}^{|V| - |H'|} \right).$$

Moreover, if $|H'|$ is greater than $2|V|/3$ and the minimum size of a hitting set in \mathcal{C} is at least $|H'| - 1$, then COMP-# d -HITTING SET can be solved in

time

$$\mathcal{O}^* \left(\binom{|H'|}{2|H'| - |V|} a_{d-1}^{|V|-|H'|} \right).$$

Proof. To prove the lemma, we give an algorithm that, for each possible partition (N, \bar{N}) of H' , computes a minimum hitting set H_N and the number h_N of minimum hitting sets subject to the constraint that these hitting sets contain all the elements of N and none of the elements of \bar{N} .

For every partition (N, \bar{N}) of H' , we either reject it as invalid or we reduce the instance (V, \mathcal{C}) to an instance (V', \mathcal{C}') by applying the following two rules in the given order.

- (H)** If there exists a set $C_i \in \mathcal{C}$ such that $C_i \subseteq \bar{N}$ then we refer to such a partition as *invalid* and reject it.
- (R)** For all sets C_i with $C_i \cap N \neq \emptyset$ put $\mathcal{C} = \mathcal{C} \setminus C_i$. In other words, all sets of \mathcal{C} , which are already hit by N , are removed.

Rule **(H)** is sound because if a set $C_i \subseteq \mathcal{C}$ is included in \bar{N} then no $\hat{H} \subseteq V \setminus \bar{N}$ contains an element of C_i and thus there is no hitting set $\hat{H} \subseteq V \setminus \bar{N}$. If a partition (N, \bar{N}) of H' is not invalid based on rule **(R)** then the instance (V, \mathcal{C}) can be reduced to the instance $I' = (V', \mathcal{C}')$, where $V' = V \setminus H'$ and $\mathcal{C}' = \{X \cap V' \mid X \in \mathcal{C} \text{ and } X \cap N = \emptyset\}$.

Summarizing, the instance I' is obtained by removing all the elements of V for which it has already been decided if they are part of H_N or not and all the sets that are hit by the elements in N . To complete H_N , it is sufficient to find a minimum hitting set of I' and to count the number of minimum hitting sets of I' . The crucial observation here is that I' is an instance of $\#\text{MHS}_{d-1}$. Indeed, H' is a hitting set of (V, \mathcal{C}) and by removing it we decrease the size of every set at least by one. Therefore, we can use an algorithm for $\#\text{MHS}_{d-1}$ to complete this step. When checking all partitions (N, \bar{N}) of H' it is straightforward to keep the accounting information necessary to compute a minimum hitting set \hat{H} and to count all minimum hitting sets.

Thus for every partition (N, \bar{N}) of H' the algorithm solving $\#\text{MHS}_{d-1}$ is called for the instance I' . There are $2^{|H'|}$ partitions (N, \bar{N}) of the vertex set H' . For each such partition, the number of elements of the instance I' is $|V'| = |V \setminus H'| = |V| - |H'|$. Thus, the running time of the algorithm is $\mathcal{O}^* \left(2^{|H'|} a_{d-1}^{|V|-|H'|} \right)$.

If $|H'| > 2|V|/3$ and the minimum size of a hitting set in \mathcal{C} is at least $|H'| - 1$, then it is not necessary to check all partitions (N, \bar{N}) of H' as the number of relevant partitions of H' becomes significantly smaller than $2^{|H'|}$, and in this case we can speed up the algorithm. Indeed, since

- $|H'| \geq |\hat{H}| \geq |H'| - 1$, and
- $|\hat{H} \cap (V \setminus H')| \leq |V| - |H'|$,

it is sufficient to consider only those partitions (N, \bar{N}) of H' such that

$$|N| \geq |H'| - 1 - (|V| - |H'|) = 2|H'| - |V| - 1.$$

In this case, the running time of the algorithm is $\mathcal{O}^* \left(\binom{|H'|}{2|H'|-|V|} a_{d-1}^{|V|-|H'|} \right)$. \square

The following lemma will be useful for the forthcoming running-time analysis.

Lemma 7. *Let n be a natural number and a be a non-negative constant. The sum of the terms $\binom{j}{2j-n} a^{n-j}$ for $j = 0, 1, \dots, n$ is upper bounded by $\mathcal{O}^* \left(\left(\frac{1 + \sqrt{1 + 4a}}{2} \right)^n \right)$.*

Proof. Let $j = n - m$. Then $\binom{j}{2j-n} a^{n-j}$ can be rewritten as $g(n, m) = \binom{n-m}{n-2m} a^m = \binom{n-m}{m} a^m$. By convention, we set $g(n, m) = 0$ whenever $n < m$. Denote by $G(n)$ the sum $\sum_{m=0}^n g(n, m)$. By a well-known decomposition of binomial coefficients, $g(n, m) = \left(\binom{n-m-1}{m-1} + \binom{n-m-1}{m} \right) a^m = a \cdot g(n-2, m-1) + g(n-1, m)$. Thus $G(n) \leq G(n-1) + a \cdot G(n-2)$. Standard calculus yields that this recurrence is asymptotically upper bounded by $\mathcal{O}^*(\alpha^n)$ where α is the largest positive root of the polynomial $x^2 - x - a = 0$, i.e. $\alpha = \frac{1 + \sqrt{1 + 4a}}{2}$. \square

Now we are ready to use iterative compression to prove the following theorem.

Theorem 8. *Suppose there exists an algorithm to solve $\#\text{MHS}_{d-1}$ in time $\mathcal{O}^*(a_{d-1}^n)$, $1 \leq a_{d-1} \leq 2$. Then $\#\text{MHS}_d$ can be solved in time*

$$\mathcal{O}^* \left(\left(\frac{1 + \sqrt{1 + 4a_{d-1}}}{2} \right)^n \right).$$

Proof. Let (V, \mathcal{C}) be an instance of $\#\text{MHS}_d$, where $V = \{v_1, v_2, \dots, v_n\}$. For $i = 1, 2, \dots, n$, let $V_i = \{v_1, v_2, \dots, v_i\}$ and $\mathcal{C}_i = \{X \in \mathcal{C} \mid X \subseteq V_i\}$. Then $I_i = (V_i, \mathcal{C}_i)$ constitutes an instance for the i^{th} stage of the iteration. We denote by H_i and h_i , a minimum hitting set of an instance I_i and the number of different minimum hitting sets of I_i respectively.

If $\{v_1\} \in \mathcal{C}$, then $H_1 = \{v_1\}$ and $h_1 = 1$; otherwise $H_1 = \emptyset$ and $h_1 = 0$.

Consider the i^{th} stage of the iteration. We have that $|H_{i-1}| \leq |H_i| \leq |H_{i-1}| + 1$ because at least $|H_{i-1}|$ elements are needed to hit all the sets of I_i except those containing element v_i and $H_{i-1} \cup \{v_i\}$ is a hitting set of I_i . Now, use Lemma 6 with $H' = H_{i-1} \cup \{v_i\}$ to compute a minimum hitting set of I_i . If $|H'| \leq 2i/3$, its running time is $\mathcal{O}^* \left(\max_{0 \leq j \leq 2i/3} \left\{ 2^j a_{d-1}^{i-j} \right\} \right) = \mathcal{O}^* \left(2^{2i/3} a_{d-1}^{i/3} \right)$ (for $a_{d-1} \leq 2$). If $|H'| > 2i/3$, the running time is

$\mathcal{O}^* \left(\max_{2i/3 < j \leq i} \left\{ \binom{j}{2j-i} a_{d-1}^{i-j} \right\} \right)$. Since for every fixed $j > 2i/3$, and $1 \leq i \leq n$,

$$\binom{j}{2j-i} a_{d-1}^{i-j} \leq \binom{j}{2j-n} a_{d-1}^{n-j},$$

the worst case running time of the algorithm is

$$\mathcal{O}^* \left(\max \left\{ \max_{1 \leq i \leq n} 2^{2i/3} a_{d-1}^{i/3}, \max_{2n/3 \leq j \leq n} \left\{ \binom{j}{2j-n} a_{d-1}^{n-j} \right\} \right\} \right).$$

Finally, $\binom{2n/3}{n/3} = 2^{2n/3}$ up to a polynomial factor, and thus the running time is $\mathcal{O}^* \left(\max_{2n/3 \leq j \leq n} \left\{ \binom{j}{2j-n} a_{d-1}^{n-j} \right\} \right)$. By Lemma 7, this latest expression is bounded by $\mathcal{O}^* \left(\left(\frac{1 + \sqrt{1 + 4a_{d-1}}}{2} \right)^n \right)$. \square

Based on the $\mathcal{O}(1.2377^n)$ algorithm for $\#MHS_2$ [26], we establish the following corollary :

Corollary 9. $\#MHS_3$ can be solved in time $\mathcal{O}(1.7198^n)$.

The same approach can be used to design an algorithm for the optimization version MHS_d , assuming that an algorithm for MHS_{d-1} is available. Based on the $\mathcal{O}(1.6278^n)$ algorithm for MHS_3 [26] this leads to an $\mathcal{O}(1.8704^n)$ time algorithm for solving MHS_4 .

Corollary 10. MHS_4 can be solved in time $\mathcal{O}(1.8704^n)$.

In the following theorem we provide an alternative approach to solve $\#MHS_d$. This is a combination of brute force enumeration (for sufficiently large hitting sets) with one application of the compression algorithm of Lemma 6. For large values of a_{d-1} , more precisely for $a_{d-1} \geq 1.6553$, this new approach gives faster algorithms than the one obtained by Theorem 8.

Theorem 11. Suppose there exists an algorithm with running time $\mathcal{O}^*(a_{d-1}^n)$, $1 < a_{d-1} \leq 2$, solving $\#MHS_{d-1}$. Then $\#MHS_d$ can be solved in time

$$\min_{0.5 \leq \alpha \leq 1} \max \left\{ \mathcal{O}^* \left(\binom{n}{\alpha n} \right), \mathcal{O}^* \left(2^{\alpha n} a_{d-1}^{n-\alpha n} \right) \right\}.$$

Proof. First the algorithm tries all subsets of V of size $\lfloor \alpha n \rfloor$ and identifies those that are a hitting set of I .

Now there are two cases. In the first case, there is no hitting set of this size. Then the algorithm verifies all sets of larger size whether they are hitting sets of I . It is straightforward to keep some accounting information to determine the number of hitting sets of the smallest size found during this enumeration phase. The running time of this phase is $\mathcal{O}^* \left(\sum_{i=\lfloor \alpha n \rfloor}^n \binom{n}{i} \right) = \mathcal{O}^* \left(\binom{n}{\alpha n} \right)$.

d	$\#\text{MHS}_d$	MHS_d
2	$\mathcal{O}(1.2377^n)$ [26]	$\mathcal{O}(1.2108^n)$ [24]
3	$\mathcal{O}(1.7198^n)$	$\mathcal{O}(1.6278^n)$ [26]
4	$\mathcal{O}(1.8997^n)$	$\mathcal{O}(1.8704^n)$
5	$\mathcal{O}(1.9594^n)$	$\mathcal{O}(1.9489^n)$
6	$\mathcal{O}(1.9824^n)$	$\mathcal{O}(1.9781^n)$
7	$\mathcal{O}(1.9920^n)$	$\mathcal{O}(1.9902^n)$

Table 1: Running times of the algorithms for $\#\text{MHS}_d$ and MHS_d .

In the second case, there exists a hitting set of size $\lfloor \alpha n \rfloor$. Then count all minimum hitting sets using the compression algorithm of Lemma 6 with H' being a hitting set of size $\lfloor \alpha n \rfloor$ found by the enumeration phase. By Lemma 6, this phase of the algorithm has running time $\mathcal{O}^*(2^{\alpha n} a_{d-1}^{n-\alpha n})$. \square

The best running times of algorithms solving $\#\text{MHS}_d$ and MHS_d are summarized in Table 1. For $\#\text{MHS}_{\geq 4}$ and $\text{MHS}_{\geq 5}$, we use the algorithm of Theorem 11. Note that the MHS_2 problem is equivalent to MVC and MIS.

Finally, it is worth to note that the technique can also be used to solve the parameterized version of d -Hitting Set. Namely, we consider the following problem:

(k, d) -HITTING SET (k -HS $_d$): Given a universe V of n elements, a collection \mathcal{C} of subsets of V of size at most d and an integer k , find a hitting set of size at most k of \mathcal{C} , if one exists.

Theorem 12. *Suppose there exists an algorithm to solve k -HS $_{d-1}$ in time $\mathcal{O}(a_{d-1}^k \cdot n^{\mathcal{O}(1)})$, where $a_{d-1} \geq 1$. Then k -HS $_d$ can be solved in time $\mathcal{O}((1 + a_{d-1})^k \cdot n^{\mathcal{O}(1)})$.*

Proof. The proof is very similar to the one of Theorem 8 except that the size of a solution is now bounded by the parameter k instead of n . Given a universe $V = \{v_1, v_2, \dots, v_n\}$ and a collection \mathcal{C} , for $i = 1, 2, \dots, n$, let $V_i = \{v_1, v_2, \dots, v_i\}$ and $\mathcal{C}_i = \{X \in \mathcal{C} \mid X \subseteq V_i\}$. Then $I_i = (V_i, \mathcal{C}_i)$ constitutes an instance of k -HS $_d$ for the i^{th} stage of the iteration. We denote by H_i a hitting set of size at most k , if one exists, of the instance I_i .

Clearly, if $\{v_1\} \in \mathcal{C}$, then $H_1 = \{v_1\}$ (assuming that $k \geq 1$); otherwise $H_1 = \emptyset$ and $h_1 = 0$.

Consider now the i^{th} stage of the iteration. The relation $|H_{i-1}| \leq |H_i| \leq |H_{i-1}| + 1 \leq k + 1$ holds since at least $|H_{i-1}|$ elements are needed to hit all the sets of I_i except those containing element v_i and $H_{i-1} \cup \{v_i\}$ is a hitting set of I_i .

d	k -HS $_d$
2	$\mathcal{O}(1.2738^k \cdot n^{\mathcal{O}(1)})$ [5]
3	$\mathcal{O}(2.0755^k \cdot n^{\mathcal{O}(1)})$ [26]
4	$\mathcal{O}(3.0755^k \cdot n^{\mathcal{O}(1)})$
5	$\mathcal{O}(4.0755^k \cdot n^{\mathcal{O}(1)})$
6	$\mathcal{O}(5.0640^k \cdot n^{\mathcal{O}(1)})$ [11]
7	$\mathcal{O}(6.0439^k \cdot n^{\mathcal{O}(1)})$ [11]

Table 2: Running times of the best known algorithms solving k -HS $_d$ for various values of d . The algorithms for $4 \leq d \leq 5$ are based on corollary 13.

As we did in Lemma 6, for every partition (N, \bar{N}) of $H_{i-1} \cup \{v_i\}$ we apply the rules **(H)** and **(R)** (see the proof of Lemma 6). Each not rejected partition (N, \bar{N}) leads to an instance $I'_i = (V'_i, \mathcal{C}'_i)$ of k -HS $_{d-1}$, where $V'_i = V_i \setminus (H_{i-1} \cup \{v_i\})$ and $\mathcal{C}'_i = \{X \cap V'_i \mid X \in \mathcal{C}_i \text{ and } X \cap N = \emptyset\}$. Thus, by using an algorithm for k -HS $_{d-1}$, a solution for of size at most k I'_i can be found, if one exists. The overall running time is given by the formula $\sum_{i=0}^k \binom{k+1}{i} a_{d-1}^{k-i} \leq 2 \cdot (1 + a_{d-1})^k$. \square

In particular, an immediate consequence of the previous theorem is the following :

Corollary 13. *Suppose there exists an algorithm to solve k -HS $_3$ in time $\mathcal{O}(a_3^k \cdot n^{\mathcal{O}(1)})$. Then, for any fixed $d \geq 4$, k -HS $_d$ can be solved in time $\mathcal{O}((a_3 + d - 3)^k \cdot n^{\mathcal{O}(1)})$.*

By using an $\mathcal{O}(2.0755^k \cdot n^{\mathcal{O}(1)})$ algorithm by Wahlström [26] for solving k -HS $_3$, we obtain the running times depicted in Table 2 for $d = 4$ and $d = 5$, which are, up to our knowledge, the best known algorithms.

4 Maximum Induced Cluster Subgraph

Clustering objects according to given similarity or distance values is an important problem in computational biology with diverse applications, e.g., in defining families of orthologous genes, or in the analysis of microarray experiments [7, 10, 14, 17, 21]. A graph theoretic formulation of the clustering problem is called CLUSTER EDITING. To define this problem we need to introduce the notion of a *cluster graph*. A graph is called a cluster graph if it is a disjoint union of cliques. In the most common parameterized version of CLUSTER EDITING the problem is stated as follows. The input is a graph $G = (V, E)$ and a positive integer integer k , the parameter. The task is to find out whether the input graph G can be transformed into a cluster

graph by adding or deleting at most k edges. This parameterized problem is fixed parameter tractable if there is an algorithm solving it in time $f(k) \cdot n^{O(1)}$, where f is an arbitrary computable function. The CLUSTER EDITING problem has been extensively studied in the realm of parameterized complexity [7, 10, 14, 21]. In this section, we study a vertex version of CLUSTER EDITING. We study the following optimization version of the problem.

MAXIMUM INDUCED CLUSTER SUBGRAPH (MICS): Given a graph $G = (V, E)$ on n vertices, find a maximum size subset $C \subseteq V$ such that $G[C]$, the subgraph of G induced by C , is a cluster graph.

Due to the following well-known observation, the MICS problem is also known as MAXIMUM INDUCED P_3 -FREE SUBGRAPH.

Observation 14. *A graph is a disjoint union of cliques if and only if it contains no induced subgraph isomorphic to the graph P_3 , the path on three vertices.*

Clearly, $C \subseteq V$ induces a cluster graph in $G = (V, E)$ (that is $G[C]$ is a disjoint union of cliques of G) iff $S = V \setminus C$ hits all induced paths on three vertices of G . Thus solving the MICS problem is equivalent to finding a minimum size set of vertices whose removal produces a maximum induced cluster subgraph of G . By Observation 14, this reduces to finding a minimum hitting set S of the collection of vertex sets of (induced) P_3 's of G . Such a hitting set S is called a P_3 -HS.

As customary when using iterative compression, we first define a compression version of the MICS problem.

COMP-MICS: Given a graph $G = (V, E)$ on n vertices and a P_3 -HS $S \subseteq V$, find a P_3 -HS of G of size at most $|S| - 1$ if one exists.

Lemma 15. *COMP-MICS can be solved in time $\mathcal{O}^*(\phi^n)$ where $\phi = (1 + \sqrt{5})/2$ is the golden ratio.*

Proof. For the proof we distinguish two cases based on the size of S .

Case 1: If $|S| \leq 2n/3$ then the following algorithm which uses matching techniques is applied.

Step 1: Enumerate all partitions of (N, \bar{N}) of S .

Step 2: For each partition, compute a maximum set $C \subseteq V$ such that $G[C]$ is a cluster graph, subject to the constraints that $N \subseteq C$ and $\bar{N} \cap C = \emptyset$, if such a set C exists.

In Step 2, we reduce the problem of finding a maximum sized C to the problem of finding a maximum weight matching in an auxiliary bipartite graph. Independent of our work, Hüffner et al. [17] also use this natural idea of reduction to weighted bipartite matching to obtain an FPT algorithm for the vertex weighted version of CLUSTER VERTEX DELETION using iterative compression. For completeness, we present the details of Step 2.

If $G[N]$ contains an induced P_3 then there is obviously no $C \subseteq V$ inducing a cluster graph that respects the partition (N, \bar{N}) . We call such a partition *invalid*.

Otherwise, $G[N]$ is a cluster graph, and thus the goal is to find a maximum size subset C' of $\bar{S} = V \setminus S$ such that $G[C' \cup N]$ is a cluster graph. Fortunately, such a set C' can be computed in polynomial time by reducing the problem to finding a maximum weight matching in an auxiliary bipartite graph.

First we describe the construction of the bipartite graph. Consider the graph $G[N \cup \bar{S}]$ and note that $G[N]$ and $G[\bar{S}]$ are cluster graphs. Now the following reduction rule is applied to the graph $G[N \cup \bar{S}]$.

(R) Remove every vertex $b \in \bar{S}$ for which $G[N \cup \{b\}]$ contains an induced P_3 .

Clearly all vertices removed by **(R)** cannot belong to any C' inducing a cluster subgraph of G . Let \hat{S} be the subset of vertices of \bar{S} which are not removed by **(R)**. Hence the current graph is $G[N \cup \hat{S}]$. Clearly $G[\hat{S}]$ is a cluster graph since $G[\bar{S}]$ is one. Further, note that no vertex of \hat{S} has neighbors in two different maximal cliques of $G[N]$ and if a vertex of \hat{S} has a neighbor in one maximal clique of $G[N]$ then it is adjacent to each vertex of this maximal clique. Thus, every vertex in \hat{S} has either no neighbor in N or it is adjacent to all the vertices of exactly one maximal clique of $G[N]$.

Now we are ready to define the auxiliary bipartite graph $G' = (A, B, E')$. Let $\{\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_r\}$ be the maximal cliques of the cluster graph $G[N]$. Let $\{\mathcal{C}'_1, \mathcal{C}'_2, \dots, \mathcal{C}'_s\}$ be the maximal cliques of the cluster graph $G[\hat{S}]$. Let $A = \{a_1, a_2, \dots, a_r, a'_1, a'_2, \dots, a'_s\}$ and $B = \{b_1, b_2, \dots, b_s\}$. Here, for all $i \in \{1, \dots, r\}$, each maximal clique \mathcal{C}_i of $G[N]$ is represented by $a_i \in A$; and for all $j \in \{1, 2, \dots, s\}$, each maximal clique \mathcal{C}'_j of $G[\hat{S}]$ is represented by $a'_j \in A$ and by $b_j \in B$.

Now there are two types of edges in G' : $a_j b_k \in E'$ if there is a vertex $u \in \mathcal{C}'_k$ such that u has a neighbor in \mathcal{C}_j , and $a'_j b_j \in E'$ if there is a vertex $u \in \mathcal{C}'_j$ such that u has no neighbor in N . Finally we define the weights for both types of edges in the bipartite graph G' . For an edge $a_j b_k \in E'$, its weight $w(a_j b_k)$ is the number of vertices in \mathcal{C}'_k being adjacent to all vertices of the maximal clique \mathcal{C}_j . For an edge $a'_j b_j$, its weight $w(a'_j b_j)$ is the number of vertices in \mathcal{C}'_j without any neighbor in N .

This transformation is of interest due to the following claim that uses the above notation.

Claim. The maximum size of a subset C' of \hat{S} such that $G[N \cup C']$ is a cluster subgraph of the graph $G^* = G[N \cup \hat{S}]$ is equal to the maximum total weight of a matching in the bipartite graph $G' = (A, B, E')$.

Proof. We first show that any matching in G' corresponds to a set $Y \subseteq \hat{S}$ that together with N induces a cluster subgraph of G^* , that is, $G[N \cup Y]$ is a P_3 -free graph. To see this, let $M = \{e_1, e_2, \dots, e_t\}$ be a matching in G' . Now if $e_l = a_j b_k$ then Y_l is the set of vertices in C'_k which are adjacent to all vertices of the maximal clique C_j . Otherwise, if $e_l = a'_j b_j$ then Y_l is the set of vertices in C'_j which have no neighbor in N . Now let us put $Y = \bigcup_{l=1}^t Y_l$. Clearly, $|Y| = \sum_{l=1}^t w(e_l)$. Now we claim that $G[N \cup Y]$ is a disjoint union of cliques. To the contrary, suppose there exists an induced P_3 in $G[N \cup Y]$, say $P = xyz$ is an induced P_3 in $G[N \cup Y]$. Then two of the vertices of P are in Y and one in N because of rule **(R)** and the fact that $G[\hat{S}]$ is a cluster graph. First let $x, z \in Y$ and $y \in N$ and $x \in C'_{t_1}$, $y \in C_{t_2}$ and $z \in C'_{t_3}$. This means selecting edges $a_{t_2} b_{t_1}$ and $a_{t_2} b_{t_3}$ in M . Secondly, let $x, y \in Y$ and $z \in N$, and thus x and y belong to the same clique C'_{t_1} , and $z \in C_{t_2}$. This means having edges $a_{t_2} b_{t_1}$ and $a'_{t_1} b_{t_1}$ in M . In both cases this contradicts M being a matching. Consequently if there is a matching M' in G^* of weight k then there is a set $Y \subseteq \hat{S}$ of size k such that $G[N \cup Y]$ is a cluster graph.

To prove the other direction, let $\{\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_q\}$ be the maximal cliques of the cluster graph $G[C']$, and let $\{\mathcal{F}'_1, \mathcal{F}'_2, \dots, \mathcal{F}'_p\}$ be the maximal cliques of the cluster graph $G[N \cup C']$. Clearly, each \mathcal{F}'_j , $1 \leq j \leq p$, contains at most one of $\{\mathcal{F}_l : 1 \leq l \leq q\}$. Let $\pi(l)$ be the integer such that $\mathcal{F}_l \subseteq \mathcal{F}'_{\pi(l)}$. If $\mathcal{F}_l = \mathcal{F}'_{\pi(l)}$ then put $e_l = a'_l b_l$. Otherwise, if $\mathcal{F}_l \subset \mathcal{F}'_{\pi(l)}$ then put $e_l = a_{\pi(l)} b_l$. Since π is injective, $M = \{e_1, e_2, \dots, e_q\}$ is a matching in G' and the definition of the weights of the edges in G' implies that the total weight of M is $\sum_{l=1}^q w(e_l) = |C'|$. Thus there is a matching of G' of total weight $|C'|$. \square

Note that the construction of the bipartite graph G' , including the application of **(R)** and the computation of a maximum weighted matching of G' can be performed in time $\mathcal{O}(n^3)$ [9]. Thus, the running time of the algorithm in Case 1 is the time needed to enumerate all subsets of S (whose size is bounded by $2n/3$) and this time is $\mathcal{O}^*(2^{2n/3}) = \mathcal{O}(1.5875^n)$.

Case 2: If $|S| > 2n/3$ then the algorithm needs to find a P_3 -HS of G of size $|S| - 1$, or show that none exists.

The algorithm proceeds as in the first case. Note that at most $n - |S|$ vertices of $V \setminus S$ can be added to N . Therefore, the algorithm verifies only those partitions (N, \bar{N}) of S satisfying $|N| \geq |S| - 1 - (n - |S|) = 2|S| - n - 1$. In this second case, the worst-case running time is

$$\mathcal{O}^* \left(\max_{2/3 < \alpha \leq 1} \left\{ \binom{\alpha n}{(2\alpha - 1)n} \right\} \right)$$

and by Lemma 7 being no more than $\mathcal{O}^*(\phi^n)$ where $\phi = (1 + \sqrt{5})/2$. \square

Now we are ready to prove the following theorem using iterative compression.

Theorem 16. *Given a graph on n vertices, MICS can be solved in time $\mathcal{O}^*(\phi^n)$ where $\phi = (1 + \sqrt{5})/2 < 1.6181$ is the golden ratio.*

Proof. Given a graph $G = (V, E)$ with vertex set $V = \{v_1, v_2, \dots, v_n\}$, let $V_i = \{v_1, v_2, \dots, v_i\}$ and let C_i be a maximum induced cluster subgraph of $G_i = G[V_i]$. Let $S_i = V_i \setminus C_i$.

The algorithm starts with $G_1, C_1 = \{v_1\}$ and $S_1 = \emptyset$. At the i^{th} iteration of the algorithm, $1 \leq i \leq n$, we maintain the invariant that we have at our disposal C_{i-1} a maximum set inducing a cluster subgraph of G_{i-1} and S_{i-1} a minimum P_3 -HS of G_{i-1} . Note that $S_{i-1} \cup \{v_i\}$ is a P_3 -HS of G_i and that no P_3 -HS of G_i has size smaller than $|S_{i-1}|$. Now use the algorithm of Lemma 15 to solve COMP-MICS on G_i with $S = S_{i-1} \cup \{v_i\}$. Then the worst-case running time is attained at the n^{th} stage of the iteration and the run time is $\mathcal{O}^*(\phi^n)$ where $\phi = (1 + \sqrt{5})/2$. \square

5 Conclusion

Iterative compression is a technique which is successfully used in the design of FPT algorithms. In this paper we show that this technique can also be used to design exact exponential time algorithms. This suggests that it might be used in other areas of algorithms as well. For example, how useful can iterative compression be in the design of approximation algorithms?

Carrying over techniques from the design of FPT algorithms to the design of exact exponential time algorithms and vice-versa is a natural and tempting idea. A challenging question in this regard is whether Measure and Conquer, a method that has been successfully used to improve the time analysis of simple exponential-time branching algorithms, can be adapted for the analysis of FPT branching algorithms.

Acknowledgements. The authors are grateful to an anonymous referee for many suggestions helping us to improve the presentation of the paper.

References

- [1] A. Björklund and T. Husfeldt, Inclusion–Exclusion Algorithms for Counting Set Partitions, in the proceedings of FOCS’06, 575–582 (2006).
- [2] A. Björklund, T. Husfeldt, P. Kaski, and M. Koivisto, Fourier meets Möbius: Fast Subset Convolution, in the proceedings of STOC’07, 67–74 (2007).

- [3] J. M. Byskov, Enumerating maximal independent sets with applications to graph colouring, *Oper. Res. Lett.* 32(6), 547–556 (2004).
- [4] J. Chen, F. V. Fomin, Y. Liu, S. Lu, and Y. Villanger, Improved Algorithms for Feedback Vertex Set Problems, *J. Comp. Syst. Sci.* 74, 2008, pp. 1188–1198.
- [5] J. Chen, I. A. Kanj, G. Xia, Improved Parameterized Upper Bounds for Vertex Cover, in the proceedings of MFCS’06, 238–249 (2006).
- [6] J. Chen, Y. Liu, S. Lu, I. Razgon, B. O’Sullivan, A Fixed-Parameter Algorithm for the Directed Feedback Vertex Set Problem, *J. ACM* 55(5), (2008).
- [7] F. Dehne, M. A. Langston, X. Luo, S. Pitre, P. Shaw and Y. Zhang, The Cluster Editing Problem: Implementations and Experiments, in the proceedings of IWPEC’06, 13-24 (2006).
- [8] F. Dehne, M. Fellows, M. Langston, F. Rosamond, and K. Stevens, An $O(2^{O(k)}n^3)$ FPT algorithm for the undirected feedback vertex set problem, *Theory Comput. Syst.* 41(3), 479–492 (2007).
- [9] J. Edmonds and R. M. Karp, Theoretical improvements in algorithmic efficiency for network flow problems, *J. ACM* 19(2), 248–264 (1972).
- [10] M. R. Fellows, M. A. Langston, F. A. Rosamond and P. Shaw, Efficient Parameterized Preprocessing for Cluster Editing, in the proceedings of FCT’07, 312-321 (2007).
- [11] H. Fernau, Parameterized Algorithms for Hitting Set: The Weighted Case, in the proceedings of CIAC’06, 332-343 (2006).
- [12] F. V. Fomin, F. Grandoni, and D. Kratsch, Measure and conquer: A simple $O(2^{0.288n})$ independent set algorithm, in the proceedings of SODA’06, 18–25 (2006).
- [13] F. V. Fomin, F. Grandoni, and D. Kratsch, A measure & conquer approach for the analysis of exact algorithms, *J. ACM* 56(5), (2009).
- [14] J. Guo, A more effective linear kernelization for cluster editing, *Theor. Comput. Sci.* 410(8-10), 718–726 (2009).
- [15] J. Guo, J. Gramm, F. Hüffner, R. Niedermeier, and S. Wernicke. Compression-based fixed-parameter algorithms for feedback vertex set and edge bipartization, *J. Comput. Syst. Sci.* 72(8), 1386–1396 (2006).
- [16] J. E. Hopcroft and R. M. Karp, An $n^{5/2}$ algorithm for maximum matching in bipartite graphs, *SIAM J. Computing* 2(4), 225–231 (1973).

- [17] F. Hüffner, C. Komusiewicz, H. Moser, and R. Niedermeier, Fixed-parameter algorithms for cluster vertex deletion, in the proceedings of LATIN'08, 711-722 (2008).
- [18] D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis, On Generating All Maximal Independent Sets, *Inf. Process. Lett.* 27(3), 119-123 (1988).
- [19] M. Koivisto, An $O(2^n)$ Algorithm for Graph Colouring and Other Partitioning Problems via Inclusion-Exclusion, in the proceedings of FOCS'06, 583-590 (2006).
- [20] J. W. Moon and L. Moser, On Cliques in Graphs, *Israel J. Mathematics* 3, 23-28 (1965).
- [21] S. Rahmann, T. Wittkop, J. Baumbach, M. Martin, A. Trub and S. Böcker, Exact and Heuristic Algorithms for Weighted Cluster Editing, in the proceedings of *Comput. Syst. Bioinformatics Conference'07*, 6 (1), 391-401 (2007).
- [22] V. Raman, S. Saurabh, and S. Sikdar, Efficient Exact Algorithms through Enumerating Maximal Independent Sets and Other Techniques, *Theory Comput. Syst.* 41(3), 1432-4350 (2007).
- [23] B. A. Reed, K. Smith, and A. Vetta, Finding odd cycle transversals, *Oper. Res. Lett.* 32 (4), 299-301 (2004).
- [24] J. M. Robson, Algorithms for maximum independent sets, *J. Algorithms* 7, 425-440 (1986).
- [25] R. Tarjan and A. Trojanowski, Finding a maximum independent set, *SIAM J. Computing* 6(3), 537-546 (1977).
- [26] M. Wahlström, Algorithms, measures and upper bounds for satisfiability and related problems, PhD thesis, Linköping University, Sweden (2007).
- [27] G. J. Woeginger, Exact algorithms for NP-hard problems: A survey, in *Combinatorial Optimization - Eureka, You Shrink!*, 185-207 (2003).