Improved Algorithms for the Feedback Vertex Set Problems [*] [**]

Jianer Chen[1], Fedor V. Fomin[2], Yang Liu[1], Songjian Lu[1], and
Yngve Villanger[2]

[1] Department of Computer Science, Texas A&M University, College Station,
Texas 77843-3112, USA
{chen,yangliu,sjlu}@cs.tamu.edu
[2] Department of Informatics, University of Bergen, N-5020 Bergen, Norway.
{fomin,yngvev}@ii.uib.no

Abstract. We present improved parameterized algorithms for the Feedback Vertex Set problem on both unweighted and weighted graphs. Both algorithms run in time $O(5^k k n^2)$. The algorithms construct a feedback vertex set of size bounded by k (in the weighted case this set is of minimum weight among the feedback vertex set of size at most k) in a given graph G, or reports that no such a feedback vertex set exists in G.

1 Introduction

Let G be a graph. A *feedback vertex set* (FVS) F in G is a set of vertices in G whose removal results in an acyclic graph (or equivalently, every cycle in G contains at least one vertex in F). The problem of finding a minimum feedback vertex set in a graph is one of the classical NP-complete problems [13] and has many applications. The history of the problem can be traced back to the early '60s. For several decades, many different algorithmic approaches were tried on this problem, including approximation algorithms, linear programming, local search, polyhedral combinatorics, and probabilistic algorithms (see the survey of Festa et al. [9]). There are also exact algorithms finding a minimum FVS in a graph on n vertices in time $O(1.9053^n)$ [16] and in time $O(1.7548^n)$ [10].

An important application of the FVS problem is *deadlock recovery* in operation systems [18], in which a deadlock is presented by a cycle in a *system resource-allocation graph* G. Therefore, in order to recover from deadlocks, we need to abort a set of processes in the system, i.e., to remove a set of vertices in the graph G, so that all cycles in G are broken. Equivalently, we need to find an FVS in G. The problem also has a version on weighted graphs, where the weight of a vertex can be interpreted as the cost of aborting the corresponding process. In this case, we are looking for an FVS in G whose weight is minimized.

In a practical system resource-allocation graph G, it can be expected that the size k of the minimum FVS in G, i.e., the number of vertices in the FVS, is fairly small. This motivated the study of *parameterized algorithms* for the FVS problem that find an FVS of k vertices in a graph of n vertices (where the weight of the FVS is minimized, in the case of weighted graphs), and run in time $f(k)n^{O(1)}$ for a fixed function f (thus, the algorithms become practically efficient when the value k is small).

This line of research has received considerable attention, most are on the FVS problem on unweighted graphs. The first group of parameterized algorithms of running time $f(k)n^{O(1)}$ for the FVS problem on unweighted graphs was given by Bodlaender [3] and by Downey and Fellows [7]. Since then a chain of dramatic improvements was obtained by different researchers (see Figure 1 for references.)

Bodlaender, Fellows [3, 7]	$O(17k^4!n^{O(1)})$
Downey and Fellows [8]	$O((2k+1)^k n^2)$
Raman et al.[15]	$O(\max\{12^k, (4\log k)^k\} n^{2.376})$
Kanj et al.[12]	$O((2\log k + 2\log\log k + 18)^k n^2)$
Raman et al.[14]	$O((12\log k/\log\log k + 6)^k n^{2.376})$
Guo et al.[11]	$O((37.7)^k n^2)$
Dehne et al.[6]	$O((10.6)^k n^3)$

Fig. 1. The history of parameterized algorithms for the unweighted FVS problem.

Our results. In this paper we use the technique of iterative compression that was already applied for several similar problems including the FVS problem [6, 11, 17]. The novel part of our approach is the new recursive procedure and its analysis which allow us to reduce the running time of the algorithm significantly. We show that the problem of finding an FVS of size k of minimum weight in a weighted graph G of n vertices can be solved in time $O(5^k kn^2)$. This improves and generalizes a long chain of results in parameterized algorithms.

We remark that randomized parameterized algorithms have also been studied in the literature for the FVS problem, for both unweighted and weighted graphs. The best known randomized parameterized algorithms for the FVS problems are due to Becker et al. [2], who developed a randomized algorithm of running time $O(4^k kn^2)$ for the FVS problem on unweighted graphs, and a randomized algorithm of running time $O(6^k kn^2)$ for the FVS problem on weighted graphs. Compared to these results, the running time of our (deterministic) algorithm comes close to that of the best randomized algorithm for the FVS problem on unweighted graphs, while our (deterministic) algorithm has been better than the previous best randomized algorithm for the FVS problem on weighted graphs.

Due to space limitations the proofs of Lemmas 4, 5 and Theorem 2 have been omitted. For a full version see [4].

2   On feedback vertex sets in unweighted graphs

In this section, we consider the FVS problem on unweighted graphs.

We start with some terminologies. A *forest* is a graph that contains no cycles. A *tree* is a forest that is connected (therefore, a forest can be equivalently defined as a collection of disjoint trees). Let W be a subset of vertices in a graph G. We will denote by $G[W]$ the subgraph of G that is induced by the vertex set W. A pair (V_1, V_2) of vertex subsets in a graph $G = (V, E)$ is a *forest bipartition* of G if $V_1 \cup V_2 = V$, $V_1 \cap V_2 = \emptyset$, and both induced subgraphs $G[V_1]$ and $G[V_2]$ are forests.

Let G be a graph and let F be a subset of vertices in G. The set F is a *feedback vertex set* (shortly, FVS) of G if $G - F$ is a forest (or equivalently, if every cycle in G contains at least one vertex in F). The *size* of an FVS F is the number of vertices in F.

Our main problem is formally defined as follows.

> FEEDBACK VERTEX SET: given a graph G and an integer k, either find an FVS of size bounded by k in G, or report that no such an FVS exists.

Before we present our algorithm for the FEEDBACK VERTEX SET problem, we first consider a special version of the problem, defined as follows:

> F-BIPARTITION FVS: given a graph G, a forest bipartition (V_1, V_2) of G, and an integer k, either find an FVS of size bounded by k for the graph G *in the subset* V_1, or report that no such an FVS exists.

Note that the main difference between the F-BIPARTITION FVS problem and the original FEEDBACK VERTEX SET problem is that we require that the FVS in the F-BIPARTITION FVS is contained in the given subset V_1.

A *bypass* operation will be used heavily in our process. Let w be a degree-2 vertex with two neighbors u and v in a graph G. We say that a graph G' is obtained from G by *bypassing* the degree-2 vertex w if G' is obtained from G by first removing the vertex w then adding a new edge between u and v.

The algorithm, **Feedback**(G, V_1, V_2, k), for the F-BIPARTITION FVS problem is given in Figure 2. We first discuss the correctness of the algorithm. The correctness of step 1 and step 2 of the algorithm is obvious. Now consider step 3. Let w be a vertex in V_1 that has at least two neighbors in V_2.

If the vertex w has two neighbors in V_2 that belong to the same tree T in the induced subgraph $G[V_2]$, then the tree T plus the vertex w contains at least one cycle. Since we are restricted to find an FVS in the vertex subset V_1, the only way to break the cycles in $T + w$ is to include the vertex w in the objective FVS. Moreover, the objective FVS of size bounded by k exists in G if and only if the remaining graph $G - w$ has an FVS of size bounded by $k - 1$ in the subset $V_1 - w$ (note that $(V_1 - w, V_2)$ is a valid forest bipartition of the graph $G - w$). Therefore, step 3.1 correctly handles this case.

If no two neighbors of the vertex w belong to the same tree in the induced subgraph $G[V_2]$, then the vertex w is either in the objective FVS or not in the

---

Algorithm-1 Feedback(G, V_1, V_2, k)
Input: $G = (V, E)$ is a graph with a forest bipartition (V_1, V_2), k is an integer
Output: An FVS F of G such that $|F| \leq k$ and $F \subseteq V_1$; or report "No" (i.e.,
            no such an FVS)

1.   **if** $(k < 0)$ or $(k = 0$ and G is not a forest) **then** return "No";
2.   **if** $(k \geq 0)$ and G is a forest **then** return $\emptyset$;
3.   **if** a vertex w in V_1 has at least two neighbors in V_2 **then**
3.1.   **if** two of the neighbors of w in V_2 belong to the same tree in $G[V_2]$
            then $F' = $ **Feedback**$(G - w, V_1 - w, V_2, k - 1)$;
                if $F' = $ "No" **then** return "No" **else** return $F' + w$;
3.2.   **else** $F_1 = $ **Feedback**$(G - w, V_1 - w, V_2, k - 1)$;
                $F_2 = $ **Feedback**$(G, V_1 - w, V_2 + w, k)$;
                if $F_1 \neq $ "No" **then** return $F_1 + w$
                else if $F_2 \neq $ "No" **then** return F_2
                else return "No";
4.   **else** pick any vertex w that has degree ≤ 1 in $G[V_1]$;
4.1.   **if** w has degree ≤ 1 in the original graph G
                then return **Feedback**$(G - w, V_1 - w, V_2, k)$
4.2.   **else** let G_w be the graph obtained from G by bypassing w;
                return **Feedback**$(G_w, V_1 - w, V_2, k)$

Fig. 2. Algorithm for unweighted FVS problem

objective FVS. If w is in the objective FVS, then we should be able to find an FVS F_1 in the graph $G - w$ such that $|F_1| \leq k - 1$ and $F_1 \subseteq V_1 - w$ (again note that $(V_1 - w, V_2)$ is a valid forest bipartition of the graph $G - w$). On the other hand, if w is not in the objective FVS, then the objective FVS for G must be contained in the subset $V_1 - w$. Also note that in this case, the subgraph $G[V_2 + w]$ induced by the subset $V_2 + w$ is still a forest since no two neighbors of w in V_2 belong to the same tree in $G[V_2]$. In consequence, $(V_1 - w, V_2 + w)$ still makes a valid forest bipartition for the graph G. Therefore, step 3.2 handles this case correctly.

Now we consider step 4. At this point, every vertex in V_1 has at most one neighbor in V_2. Moreover, since the induced subgraph $G[V_1]$ is a forest, there must be a vertex w in V_1 that has degree bounded by 1 in $G[V_1]$ (note that V_1 cannot be empty at this point since otherwise the algorithm would have stopped at step 2). If the vertex w also has degree bounded by 1 in the original graph G, then removing w does not help breaking any cycles in G. Therefore, the vertex w can be discarded. This case is correctly handled by step 4.1. Otherwise, the vertex w has degree bounded by 1 in the induced subgraph $G[V_1]$ but has degree larger than 1 in the original graph G. Observing also the fact that w has at most one neighbor in V_2, we can easily derive in this case that the degree of w in the original graph G must be 2, and that w has two neighbors u_1 and u_2 such that u_1 is in V_1 and u_2 is in V_2. Therefore, if w is in the objective FVS F, then the set $F' = F - w + u_1$ will also make a valid solution to the given problem instance. Thus, by bypassing the degree-2 vertex w in G, we obtain a graph G_w, with the

forest bipartition $(V_1 - w, V_2)$, such that G_w has an FVS of size bounded by k in $V_1 - w$ if and only if the original graph G has an FVS of size bounded by k in V_1. In conclusion, step 4.2 correctly handles this case.

Now we are ready to present the following lemma.

Lemma 1. *The algorithm* **Feedback**(G, V_1, V_2, k) *solves the* F-BIPARTITION FVS *problem correctly. The running time of the algorithm is bounded by* $O(2^{k+l}n^2)$, *where n is the number of vertices in the graph G, and l is the number of connected components (i.e. trees) in the induced subgraph* $G[V_2]$.

Proof. The correctness of the algorithm has been verified by the above discussion. Now we consider the complexity of the algorithm.

The recursive execution of the algorithm can be described as a search tree $\mathcal{T}$. We first count the number of leaves in the search tree $\mathcal{T}$. Note that only step 3.2 of the algorithm corresponds to branches in the search tree $\mathcal{T}$. Let $T(k, l)$ be the total number of leaves in the search tree $\mathcal{T}$ for the algorithm **Feedback**(G, V_1, V_2, k), where l is the number of trees in the forest $G[V_2]$. Inductively, the number of leaves in the search tree $\mathcal{T}_1$ corresponding to the recursive call **Feedback**$(G - w, V_1 - w, V_2, k - 1)$ is bounded by $T(k - 1, l)$. Moreover, we assume at step 3.2 that w has at least two neighbors in V_2 and that no two neighbors of w in V_2 belong to the same tree in $G[V_2]$. Therefore, the vertex w "merges" at least two trees in $G[V_2]$ into a single tree in $G[V_2 + w]$. Hence, the number of trees in $G[V_2 + w]$ is bounded by $l - 1$. In consequence, the number of leaves in the search tree $\mathcal{T}_2$ corresponding to the recursive call **Feedback**$(G, V_1 - w, V_2 + w, k)$ is bounded by $T(k, l - 1)$. This gives the following recurrence relation:

$$T(k, l) \leq T(k - 1, l) + T(k, l - 1)$$

It is easy to derive from this relation that $T(k, l) = O(2^{k+l})$. Finally, observe that along each root-leaf path in the search tree $\mathcal{T}$, the total number of executions of steps 1, 2, 3.1, and 4 of the algorithm is bounded by $O(n)$ because each of these steps either stops immediately, or reduces the input graph size by at least 1. Moreover, it is also easy to verify that each of these steps takes time $O(n)$.

Therefore, the computation time along each root-leaf path in the search tree $\mathcal{T}$ is bounded by $O(n^2)$. In conclusion, the running time of the algorithm **Feedback**(G, V_1, V_2, k) is bounded by $O(2^{k+l}n^2)$. This completes the proof of the theorem. $\square$

Following the idea of *iterative compression* proposed by Reed et al. [17], we formulate the following problem:

FVS REDUCTION: given a graph G and an FVS F of size $k + 1$ for G, either construct an FVS of size bounded by k for G, or report that no such an FVS exists.

Lemma 2. *The* FVS REDUCTION *problem can be solved in time* $O(5^k n^2)$.

Proof. We use the algorithm **Feedback** to solve the FVS REDUCTION problem. Let F be the FVS of size $k+1$ in the graph $G = (V, E)$. Every FVS F' of size bounded by k for G is a union of a subset F_1 of at most $k - j$ vertices in $V - F$ and a subset F_2 of j vertices in F, for some integer j, $0 \leq j \leq k$. Note that since we assume that no vertex in $F - F_2$ is in the FVS F', the induced subgraph $G[F - F_2]$ must be a forest. Therefore, for each j, $0 \leq j \leq k$, we enumerate all subsets of j vertices in F. For each such a subset F_2 in F such that $G[F - F_2]$ is a forest, we seek a subset F_1 of at most $k - j$ vertices in $V - F$ such that $F_1 \cup F_2$ makes an FVS for the graph G.

Fix a subset F_2 in F, where $|F_2| = j$. Note that the graph G has an FVS $F_1 \cup F_2$ of size bounded by k, where $F_1 \subseteq V - F$, if and only if the subset F_1 of $V - F$ is an FVS for the graph $G - F_2$ and the size of F_1 is bounded by $k - j$. Therefore, to solve the original problem, we can instead consider how to construct an FVS F_1 for the graph $G - F_2$ such that $|F_1| \leq k - j$ and $F_1 \subseteq V - F$.

Since F is an FVS for G, we have that the induced subgraph $G[V - F] = G - F$ is a forest. Moreover, by our assumption, the induced subgraph $G[F - F_2]$ is also a forest. Note that $(V - F) \cup (F - F_2) = V - F_2$, which is the vertex set for the graph $G' = G - F_2$. Therefore, $(V - F, F - F_2)$ is a forest bipartition of the graph G'. Thus, an FVS F_1 for the graph G' such that $|F_1| \leq k - j$ and $F_1 \subseteq V - F$ can be constructed by the algorithm **Feedback**$(G', V - F, F - F_2, k - j)$.

Since $|F| = k+1$ and $|F_2| = j$, we have that $|F - F_2| = k + 1 - j$. Therefore, the forest $G[F - F_2]$ contains at most $k + 1 - j$ trees. By Lemma 1, the running time of the algorithm **Feedback**$(G', V - F, F - F_2, k - j)$ is bounded by $O(2^{(k-j)+(k+1-j)} n^2) = O(4^{k-j} n^2)$. Now for all integers j, $0 \leq j \leq k$, we enumerate all subsets F_2 of j vertices in F and apply the algorithm **Feedback**$(G', V - F, F - F_2, k - j)$ for those F_2 such that $G[F - F_2]$ is a forest. As we discussed above, the graph G has an FVS of size bounded by k if and only if for some $F_2 \subseteq F$, the algorithm **Feedback**$(G', V - F, F - F_2, k - j)$ produces an FVS F_1 for the graph G'. The running time of this process is bounded by

$$\sum_{j=0}^{k} \binom{k+1}{j} \cdot O(4^{k-j} n^2) = \sum_{j=0}^{k} \binom{k+1}{k-j+1} O(4^{k-j+1} n^2) = O(5^k n^2).$$

This completes the proof of the lemma.   □

Finally, by combining Lemma 2 with iterative compression, we obtain the main result of this section.

Theorem 1. *The* FEEDBACK VERTEX SET *problem is solvable in time* $O(5^k k n^2)$.

Proof. To solve the FEEDBACK VERTEX SET problem, for a given graph $G = (V, E)$, we start by applying Bafna et al.'s 2-approximation algorithm for the MINIMUM FEEDBACK VERTEX SET problem [1]. This algorithm runs in $O(n^2)$ time, and either returns an FVS F' of size bounded by $2k$, or verifies that no FVS of size bounded by k exists. If no FVS is returned, the algorithm is terminated with the result that no FVS of size bounded by k exists. In the

case of the opposite result, we use any subset $V' \subseteq F'$ of k vertices, and let $V_0 = V' \cup (V - F')$. Of course, the induced subgraph $G[V_0]$ has an FVS of size k, since $G[V_0 - V']$ is a forest. Let $F' - V_0 = \{v_1, v_2, \ldots, v_{|F'|-k}\}$, and let $V_i = V_0 \cup \{v_1, \ldots, v_i\}$ for $i = 0, 1, \ldots, |F'| - k$. Inductively, suppose that we have constructed an FVS F_i for the graph $G[V_i]$, where $|F_i| = k$. Then the set $F'_{i+1} = F_i + v_{i+1}$ is obviously an FVS for the graph $G[V_{i+1}]$ and $|F'_{i+1}| = k + 1$.

Now the pair $(G[V_{i+1}], F'_{i+1})$ is an instance for the FVS REDUCTION problem. Therefore, in time $O(5^k n^2)$, we can either construct an FVS F_{i+1} of size k for the graph $G[V_{i+1}]$, or report no such an FVS exists. Note that if the graph $G[V_{i+1}]$ does not have an FVS of size k, then the original graph G cannot have an FVS of size k. In this case, we simply stop and claim the non-existence of an FVS of size k for the original graph G. On the other hand, with an FVS F_{i+1} of size k for the graph $G[V_{i+1}]$, our induction proceeds to the next graph $G[V_{i+1}]$, until we reach the graph $G = G[V_{|F'|-k}]$. Clearly, this process runs in time $O(5^k k n^2)$ since $|F'| - k \leq k$, and solves the FEEDBACK VERTEX SET problem. $\square$

3   On feedback vertex sets in weighted graphs

In this section, we discuss the FVS problem on weighted graphs. A weighted graph $G = (V, E)$ is an undirected graph, where each vertex $u \in V$ is assigned a non-negative weight. The weight of a vertex set $A \subseteq V$ is the sum of the vertex weights of all vertices in A. We denote by $|A|$ the size of, i.e., the number of vertices in, the set A . The (parameterized) feedback vertex set problem on weighted graphs is formally defined as follows:

> WEIGHTED-FVS: given a weighted graph G and an integer k, either find an FVS F of minimum weight for G such that $|F| \leq k$, or report that no FVS of size bounded by k exists in G.

Our algorithm for the weighted case has several similarities with the unweighted case, but also has a significant difference. The difference is that the bypass operation for unweighted graphs can no longer be used in the weighted case. Indeed, a degree-2 vertex in a weighted graph may be necessarily included in the objective FVS if its weight is very small.

On the other hand, if two degree-2 vertices v and w are adjacent, then we can always bypass the one with a larger weight. This is because every cycle in the graph either contains both v and w or contains neither of them, so we can always assume that the one with a larger weight is not included in the objective FVS. We call this operation that bypasses the vertex of a larger weight in two adjacent degree-2 vertices in a weighted graph the *restricted bypass* operation.

However, since the restricted bypass operation cannot guarantee to eliminate all degree-2 vertices in a weighted graph, step 4.2 in the algorithm **Feedback** is not always possible. To overcome this difficulty, we introduce a new partition structure of the vertices in a weighted graph.

A triple (V_0, V_1, V_2) is a *independent-forest partition* (briefly, an *IF-partition*) of a graph $G = (V, E)$ if (V_0, V_1, V_2) is a partitioning of V such that (1) $G[V_1]$ and

$G[V_2]$ are forests; (2) $G[V_0]$ is an independent set; and (3) every vertex $u \in V_0$ is of degree 2 in G, and all neighbors of u are in V_2.

We consider the following problem on weighted graphs that is similar to the F-BIPARTITION FVS problem on unweighted graphs.

> WEIGHTED IF-PARTITION FVS: given a weighted graph G, an IF-partition (V_0, V_1, V_2) of G, and an integer k, either find an FVS F of minimum weight for G that satisfies the conditions $|F| \leq k$ and $F \subseteq V_0 \cup V_1$, or report that no such an FVS exists.

To study WEIGHTED IF-PARTITION FVS, we introduce the following concept.

Definition 1. Let (G, k) be an instance of WEIGHTED IF-PARTITION FVS, where an IF-partition (V_0, V_1, V_2) of G is given. The *deficiency* of (G, k) is defined by

$$\tau(k, V_0, V_1, V_2) = k - (|V_0| - \#c(V_2) + 1)$$

where $\#c(V_2)$ is the number of connected components in the subgraph $G[V_2]$.

Intuitively, the deficiency $\tau(k, V_0, V_1, V_2)$ of the instance (G, k) is the maximum number of vertices in the objective FVS that are in the set V_1 (this will become clearer during our discussion below). Our algorithm for the WEIGHTED IF-PARTITION FVS problem is based on the following observation: once we have correctly determined all vertices in the objective FVS that are in the set V_1, the problem will become solvable in polynomial time, as shown in the following lemma.

Lemma 3. *Let (G, k) be an instance of* WEIGHTED IF-PARTITION FVS*, with an IF-partition (V_0, V_1, V_2) of G. If $V_1 = \emptyset$ or $\tau(k, V_0, V_1, V_2) \leq 0$, then the solution to the instance (G, k) can be constructed in time $O(n^2)$.*

Proof. Construct a new graph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$, where each vertex μ in $\mathcal{V}$ corresponds to a connected component in the induced subgraph $G[V_2]$, and each edge $[\mu, \nu]$ in $\mathcal{E}$ corresponds to a vertex v in the set V_0 such that the two edges incident to v in G are connected to the connected components in $G[V_2]$ that correspond to the two vertices μ and ν, respectively, in $\mathcal{H}$. Equivalently, the graph $\mathcal{H}$ can be obtained from the induced subgraph $G[V_0 \cup V_2]$ by shrinking each connected component in $G[V_2]$ into a single vertex and bypassing each degree-2 vertex in V_0. Moreover, we give each edge in $\mathcal{H}$ a weight that is equal to the weight of the corresponding vertex in V_0. Thus, the graph $\mathcal{H}$ is a graph with edge weights.

First consider the case of $V_1 = \emptyset$. If $k < 0$ then the solution to (G, k) is "No": we cannot remove a negative number of vertices. Assuming $k \geq 0$. Then we need to find a minimum-weight subset of at most k vertices in the set V_0 whose removal from $G = G[V_0 \cup V_2]$ results in an acyclic graph. Note that removing vertices in V_0 in the graph G corresponds to removing edges in the graph $\mathcal{H}$. Therefore, this problem is equivalent to finding a minimum-weight subset of at most k edges in the graph $\mathcal{H}$ whose removal from $\mathcal{H}$ results in an acyclic graph (note that each connected component in $G[V_2]$ is a tree). Let $\mathcal{H}_1$,

..., $\mathcal{H}_s$ be the connected components of the graph $\mathcal{H}$, where for each i, the component $\mathcal{H}_i$ has n_i vertices and m_i edges. In order to get an acyclic graph from $\mathcal{H}$, it is necessary and sufficient to remove $m_i - n_i + 1$ edges from $\mathcal{H}_i$ for each i (i.e., to make each connected component in $\mathcal{H}$ a tree). In consequence, in order to get an acyclic graph from $\mathcal{H}$, it is necessary and sufficient to remove $\sum_i^s (m_i - n_i + 1) = |\mathcal{E}| - |\mathcal{V}| + s$ edges from the graph $\mathcal{H}$.

Correspondingly, in case $V_1 = \emptyset$, a minimum-weight FVS in V_0 for the graph G contains exact $|\mathcal{E}| - |\mathcal{V}| + s$ vertices. Note that $|\mathcal{E}| = |V_0|$, and $|\mathcal{V}|$ is equal to the number $\#c(V_2)$ of connected components in the induced subgraph $G[V_2]$. Thus, every FVS in the graph G contains at least $|V_0| - \#c(V_2) + s$ vertices. Therefore, if $\tau(k, V_0, V_1, V_2) = k - (|V_0| - \#c(V_2) + 1) < 0$, or $\tau(k, V_0, V_1, V_2) = 0$ but $s > 1$, then we have $k < |V_0| - \#c(V_2) + s$. That is, the graph G has no FVS of size bounded by k and the solution to the instance (G, k) is a "No".

The remaining case is that $s = 1$, and $\tau(k, V_0, V_1, V_2) = k - (|V_0| - \#c(V_2) + 1) = 0$. In this case, to find a minimum-weight FVS of k vertices in V_0, we construct a maximum-weight spanning tree in the graph $\mathcal{H}$ (this can be done by a modified minimum spanning tree algorithm of time $O(n^2)$ [5]). The remaining $|\mathcal{E}| - |\mathcal{V}| + 1 = |V_0| - \#c(V_2) + 1 = k$ edges in $\mathcal{H}$ then correspond to k vertices in the set V_0 that make a minimum-weight FVS for the graph G. Summarizing the above discussion, we conclude that if $V_1 = \emptyset$, then the solution to the instance (G, k) can be constructed in time $O(n^2)$.

Now consider the case $\tau(k, V_0, V_1, V_2) \leq 0$. As shown above, even to break all cycles in the induced subgraph $G[V_0 \cup V_2]$ requires removing at least $|V_0| - \#c(V_2) + 1$ vertices in the set V_0. Therefore, if $\tau(k, V_0, V_1, V_2) \leq 0$, then $k \leq (|V_0| - \#c(V_2) + 1)$, and all k vertices in the objective FVS must be in the set V_0 in order to break all cycles in the induced subgraph $G[V_0 \cup V_2]$, and no vertex in the objective FVS can be in the set V_1. Hence, if the induced subgraph $G[V_1 \cup V_2]$ contains a cycle, then the solution to (G, k) is a "No". On the other hand, suppose that $G[V_1 \cup V_2]$ is a forest, then the graph G has another IF-partition (V_0', V_1', V_2'), where $V_0' = V_0$, $V_1' = \emptyset$, and $V_2' = V_1 \cup V_2$. It is easy to verify that in this case the instance (G, k) with the IF-partition (V_0', V_1', V_2') has the same solution set as the same instance with the IF-partition (V_0, V_1, V_2). Now since $V_1' = \emptyset$, by the first part of this lemma, the solution to (G, k) with the IF-partition (V_0', V_1', V_2') can be constructed in time $O(n^2)$. This completes the proof of the lemma.   □

Now we are ready for our main algorithm, which is given in Figure 3 and solves the WEIGHTED IF-PARTITION FVS problem. As explained for the unweighted case, vertices of degree less than 2 cannot contribute to the objective FVS, thus can be directly deleted. Moreover, each restricted bypass operation takes time $O(n)$ and eliminates a degree-2 vertex in a pair of adjacent degree-2 vertices. Therefore, we can perform a preprocessing of time $O(n^2)$ and assume that the input graph G of the algorithm contains no vertex of degree less than 2, and no two adjacent degree-2 vertices. Moreover, for each tree in the forest $G[V_1]$, we fix a root so that we can talk about the "lowest leaf" in a tree in $G[V_1]$.

Algorithm-1 W-Feedback(G, V_0, V_1, V_2, k)
Input: $G = (V, E)$ is a graph with an IF-partition (V_0, V_1, V_2), k is an integer
Output: a minimum-weight FVS F of G such that $|F| \leq k$ and $F \subseteq V_0 \cup V_1$;
          or report "No" (i.e., no such an FVS).
1. **if** $V_1 = \emptyset$ or $\tau(k, V_0, V_1, V_2) \leq 0$ **then** solve the problem in time $O(n^2)$;
2. **if** $(k < 0)$ or $(k = 0$ and G is not a forest) **then** return "No";
3. **if** a vertex w in V_1 is incident to 2 edges whose other ends are in V_2 **then**
3.1   **if** 2 edges incident to w have their other ends in the same tree in $G[V_2]$
            then return $w + $ **W-Feedback**$(G - w, V_0, V_1 - w, V_2, k - 1)$;
3.2   **else** $F_1 = w + $ **W-Feedback**$(G - w, V_0, V_1 - w, V_2, k - 1)$;
            $F_2 = $ **W-Feedback**$(G, V_0, V_1 - w, V_2 + w, k)$;
            return the one of F_1 and F_2 that has a smaller weight;
4. **else** pick a lowest leaf w_1 in any tree T in $G[V_1]$;
4.1   let w be the parent of w_1 in T, and let $w_1, \ldots, w_t$ be the children
            of w in T, where for each i, w_i has a neighbor v_i in V_2;
4.2   **if** w has a neighbor v in V_2 **then**
4.2.1   **if** for some i, v and v_i are in the same tree in $G[V_2]$
4.2.2   **then** $F_1 = w + $ **W-Feedback**$(G - w, V_0, V_1 - w, V_2, k - 1)$;
                $F_2 = w_i + $ **W-Feedback**$(G - w_i, V_0, V_1 - w_i, V_2, k - 1)$;
                return the one of F_1 and F_2 that has a smaller weight;
4.2.3   **else** $F_1 = w + $ **W-Feedback**$(G - w, V_0, V_1 - w, V_2, k - 1)$;
4.2.4 $F_2 = $ **W-Feedback**$(G, V_0 + w_1 + \cdots + w_t, V_1 - w, V_2 + w, k)$;
                return the one of F_1 and F_2 that has a smaller weight;
4.3   **else** $F_1 = w + $ **W-Feedback**$(G - w, V_0, V_1 - w, V_2, k - 1)$;
            $F_2 = $ **W-Feedback**$(G, V_0 + w_1 + \cdots + w_t, V_1 - w, V_2 + w, k)$;
            return the one of F_1 and F_2 that has a smaller weight;

Fig. 3. Algorithm for weighted FVS problem

We first discuss the correctness of the algorithm. Step 1 of the algorithm is justified by Lemma 3. Justifications for steps 2, 3.1, and 3.2 are exactly the same as those for the unweighted case. Now consider step 4. When the algorithm reaches step 4, the following conditions hold: (1) every vertex in G has degree at least 2; (2) there are no two adjacent degree-2 vertices in G; (3) $V_1 \neq \emptyset$; and (4) every vertex in V_1 is incident to at most one edge whose other end is in V_2. Conditions (1) and (2) hold because of our assumption on the input graph G; condition (3) holds because of step 1; and condition (4) holds because of step 3.

By condition (3) and because the induced subgraph $G[V_1]$ is a forest, step 4 can always pick the vertex w_1. By conditions (1) and (4), the vertex w_1 is adjacent to a unique vertex v_1 in V_2. Then by condition (1) again, w_1 must have a parent w in the tree T in $G[V_1]$. In consequence, the vertex w_1 has degree exactly 2 in the graph G. Finally, since w_1 is the lowest leaf in the tree T, all children $w_1, \ldots, w_t$ of w in T are also leaves in T. By conditions (1) and (4) again, each child w_i of w has a unique neighbor v_i in the set V_2, and every child w_i of w has degree exactly 2 in the graph G.

If the vertex w has a (unique) neighbor v in V_2, and the vertices v and v_i for some i belong to the same tree T' in $G[V_2]$, as given in step 4.2.1, then this tree T' plus the edges $[v_i, w_i]$, $[w_i, w]$, and $[w, v]$ must contain a cycle, and in this cycle all vertices are in V_2 except w_i and w. Therefore, to break this cycle using vertices not in V_2, one of the vertices w and w_i must be included in the objective FVS. Thus, step 4.2.2 correctly handles this case. On the other hand, suppose that the vertex v is in a tree in $G[V_2]$ that does not contain any of the vertices $v_1, \ldots, v_t$. Then we simply branch on the vertex w: step 4.2.3 includes w in the objective FVS, and step 4.2.4 excludes w from the objective FVS, by moving w from V_1 to V_2. Note that in case of step 4.2.4, each of the degree-2 vertices $w_1, \ldots, w_t$ now is incident to two edges whose other ends are in V_2. Thus, we can correctly move these vertices from V_1 to V_0. Moreover, it is easy to verify that the new partition in each of the cases still makes a valid IF-partition of the graph G.

Finally, suppose that the vertex w has no neighbor in V_2. By condition (2) and because the vertex w_1 has degree 2 in G, the vertex w must have at least two children in the tree T (i.e., $t \geq 2$). Again we branch on w by either including or excluding w in the objective FVS, as given by step 4.3.

Since all possible cases are covered in the algorithm, we conclude that when the algorithm **W-Feedback** stops, it must output a correct solution to the given instance (G, k).

Lemma 4. *The algorithm* **W-Feedback**(G, V_0, V_1, V_2, k) *solves the* WEIGHTED IF-BIPARTITION FVS *problem correctly, and runs in time* $O(2^{\tau(k, V_0, V_1, V_2)} n^2)$, *where n is the number of vertices in the graph G.*

With Lemma 4, we can now proceed the same way as for the unweighted case to solve the original WEIGHTED-FVS problem. Consider the following weighted version of the FVS REDUCTION problem.

> WEIGHTED FVS REDUCTION: given a weighted graph G and an FVS F of size $k + 1$ for G, either construct an FVS F' of minimum weight that satisfies $|F'| \leq k$, or report that no such an FVS exists.

Note that in the definition of WEIGHTED FVS REDUCTION, we do not require that the given FVS F' of size $k + 1$ have the minimum weight.

Lemma 5. *The* WEIGHTED FVS REDUCTION *problem is solvable in time* $O(5^k n^2)$.

Now using Theorem 1 and Lemma 5, we obtain the main result of this paper.

Theorem 2. *The* WEIGHTED-FVS *problem is solvable in time* $O(5^k k n^2)$.

Acknowledgment:

References

1. V. Bafna, P. Berman, and T. Fujito. A 2-approximation algorithm for the undirected feedback vertex set problem. *SIAM J. Discrete Math.*, 12(3):289–297, 1999.
2. A. Becker, R. Bar-Yehuda, and D. Geiger. Randomized algorithms for the loop cutset problem. *J. Artif. Intell. Res. (JAIR)*, 12:219–234, 2000.
3. H. L. Bodlaender. On disjoint cycles. *Int. J. Found. Comput. Sci.*, 5(1):59–68, 1994.
4. J. Chen, F.V. Fomin, Y. Liu, S. Lu, and Y. Villanger. Improved algorithms for the feedback vertex set problems. Technical Report 348, University of Bergen, Norway, 2007.
5. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, Second Edition*. The MIT Press and McGraw-Hill Book Company, 2001.
6. F. K. H. A. Dehne, M. R. Fellows, M. A. Langston, F. A. Rosamond, and K. Stevens. An o($2^{o(k)}n^3$) fpt algorithm for the undirected feedback vertex set problem. *COCOON*, volume 3595 of *Lecture Notes in Computer Science*, pages 859–869. Springer, 2005.
7. R. G. Downey and M. R. Fellows. Fixed Parameter Tractability and Completeness. *Complexity Theory: Current Research*, pages 191–225. Cambridge University Press, 1992.
8. R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Springer-Verlag, New York, 1999.
9. Paola Festa, Panos M. Pardalos, and Mauricio G. C. Resende. Feedback set problems. In *Handbook of combinatorial optimization, Supplement Vol. A*, pages 209–258. Kluwer Acad. Publ., Dordrecht, 1999.
10. F. V. Fomin, S. Gaspers, and A. V. Pyatkin. Finding a minimum feedback vertex set in time $O(1.7548^n)$. *IWPEC*, volume 4169 of *Lecture Notes in Computer Science*, pages 184–191. Springer, 2006.
11. 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.
12. I. A. Kanj, M. J. Pelsmajer, and M. Schaefer. Parameterized algorithms for feedback vertex set. *IWPEC*, volume 3162 of *Lecture Notes in Computer Science*, pages 235–247. Springer, 2004.
13. R. M. Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Plenum Press, New York, 1972.
14. V. Raman, S. Saurabh, and C. R. Subramanian. Faster fixed parameter tractable algorithms for finding feedback vertex sets. *ACM Trans. Algorithms*, 2(3):403–415, 2006.
15. V. Raman, S. Saurabh, and C. R. Subramanian. Faster fixed parameter tractable algorithms for undirected feedback vertex set. *ISAAC 2002*, volume 2518 of *Lecture Notes in Computer Science*, pages 241–248. Springer, 2002.
16. I. Razgon. Exact computation of maximum induced forest. *SWAT*, volume 4059 of *Lecture Notes in Computer Science*, pages 160–171. Springer, 2006.
17. B. Reed, K. Smith, and A. Vetta. Finding odd cycle transversals. *Oper. Res. Lett.*, 32(4):299–301, 2004.
18. A. Silberschatz and P. Galvin. *Operating System Concepts, 4th edition*. Addison-Wesley, 1994.