



The Power of Linear-Time Data Reduction for Maximum Matching

George B. Mertzios¹ · André Nichterlein² · Rolf Niedermeier²

Received: 18 January 2019 / Accepted: 13 June 2020 / Published online: 6 July 2020
© The Author(s) 2020

Abstract

Finding maximum-cardinality matchings in undirected graphs is arguably one of the most central graph primitives. For m -edge and n -vertex graphs, it is well-known to be solvable in $O(m\sqrt{n})$ time; however, for several applications this running time is still too slow. We investigate how linear-time (and almost linear-time) data reduction (used as preprocessing) can alleviate the situation. More specifically, we focus on *linear-time kernelization*. We start a deeper and systematic study both for general graphs and for bipartite graphs. Our data reduction algorithms easily comply (in form of preprocessing) with every solution strategy (exact, approximate, heuristic), thus making them attractive in various settings.

Keywords Maximum-cardinality matching · Bipartite graphs · Linear-time algorithms · Kernelization · Parameterized complexity analysis · FPT in P

A short version of this article appeared in the *Proceedings of the 42nd International Symposium on Mathematical Foundations of Computer Science (MFCS '17)*, pages 46:1–46:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. This article now contains all proofs in full detail. Supported by the EPSRC grant EP/P020372/1, by the DFG project FPTinP, NI 369/16, and by a postdoc fellowship of the German Academic Exchange Service (DAAD) while the second author was at Durham University.

✉ André Nichterlein
andre.nichterlein@tu-berlin.de

George B. Mertzios
george.mertzios@durham.ac.uk

Rolf Niedermeier
rolf.niedermeier@tu-berlin.de

¹ Department of Computer Science, Durham University, Durham, UK

² Algorithmics and Computational Complexity, Faculty IV, TU Berlin, Berlin, Germany

1 Introduction

“Matching is a powerful piece of algorithmic magic” [24]. In the MAXIMUM MATCHING problem, given an undirected graph, one has to compute a maximum-cardinality set of nonoverlapping edges. Maximum matching is arguably among the most fundamental graph-algorithmic primitives allowing for a polynomial-time algorithm. More specifically, on an n -vertex and m -edge graph a maximum matching can be found in $O(m\sqrt{n})$ time [21]. Improving this upper time bound resisted decades of research. Recently, however, Duan and Pettie [8] presented a linear-time algorithm that computes a $(1 - \epsilon)$ -approximate maximum-weight matching, where the running time dependency on ϵ is $\epsilon^{-1} \log(\epsilon^{-1})$. For the unweighted case, the $O(m\sqrt{n})$ algorithm of Micali and Vazirani [21] implies a linear-time $(1 - \epsilon)$ -approximation, where in this case the running time dependency on ϵ is ϵ^{-1} [8]. We take a different route: First, we do not give up the quest for optimal solutions. Second, we focus on efficient—more specifically, linear-time executable—data reduction rules, that is, not solving an instance but significantly shrinking its size before actually solving the problem.¹ In the context of decision problems and parameterized algorithmics this approach is known as kernelization; this is a particularly active area of algorithmic research on NP-hard problems.

The spirit behind our approach is thus closer to the identification of efficiently solvable special cases of MAXIMUM MATCHING. There is quite some body of work in this direction. For instance, since an augmenting path can be found in linear time [10], the standard augmenting path-based algorithm runs in $O(s(n + m))$ time, where s is the number of edges in the maximum matching. Yuster [26] developed an $O(rn^2 \log n)$ -time algorithm, where r is the difference between maximum and minimum degree of the input graph. Moreover, there are linear-time algorithms for computing maximum matchings in special graph classes, including convex bipartite [25], strongly chordal [7], chordal bipartite graphs [6], and cocomparability graphs [20].

All this and the more general spirit of “parameterization for polynomial-time solvable problems” (also referred to as “FPT in P” or “FPTP” for short) [12] forms the starting point of our research. Remarkably, Fomin et al. [9] recently developed an algorithm to compute a maximum matching in graphs of treewidth k in $O(k^4 n \log^2 n)$ randomized time. Afterwards, Iwata, Ogasawara, and Ohsaka [16] provided an elegant algorithm computing a maximum matching in graphs of treedepth ℓ in $O(\ell \cdot m)$ time. This implies an $O(k^2 n \log n)$ -time algorithm where k is the treewidth, since $m \in O(kn)$ and $\ell \leq (k + 1) \log n$ [23]. Recently, Kratsch and Nelles [19] presented a $O(r^2 \log(r)n + m)$ -time algorithm where r is the modular-width.

Following the paradigm of *kernelization*, that is, provably effective and efficient data reduction, we provide a systematic exploration of the power of not only polynomial-time but actually linear-time data reduction for MAXIMUM MATCHING. Thus, our aim (fitting within FPTP) is to devise problem kernels that are computable in

¹ Doing so, however, we focus on the unweighted case.

Table 1 Our kernelization results

Parameter k	Running time	Kernel size	
Results for MATCHING			
Feedback edge number	$O(n + m)$ time	$O(k)$ vertices and edges	(Theorem 1)
Feedback vertex number	$O(kn)$ time	$2^{O(k)}$ vertices and edges	(Theorem 2)
Results for BIPARTITE MATCHING			
Distance to chain graphs	$O(n + m)$ time	$O(k^3)$ vertices	(Theorem 3)

linear time. In other words, the fundamental question we pose is whether there is a very efficient preprocessing that provably shrinks the input instance, where the effectiveness is measured by employing some parameters. The philosophy behind is that if we can design linear-time data reduction algorithms, then we may employ them for free before afterwards employing any super-linear-time solving algorithm. We believe that this sort of question deserves deeper investigation and we initiate it based on the MAXIMUM MATCHING problem. In fact, in follow-up work we demonstrated that such linear-time data reduction rules can significantly speed-up state-of-the-art solvers for MATCHING [18].

As kernelization is usually defined for decision problems, we use in the remainder of the paper the *decision version* of MAXIMUM MATCHING. In the rest of the paper we call this decision version MATCHING. In a nutshell, a kernelization of a decision problem instance is an algorithm that produces an equivalent instance whose size can solely be upper-bounded by a function in the parameter (preferably a polynomial function). The focus on decision problems is justified by the fact that all our results, although formulated for the decision version, in a straightforward way extend to the corresponding optimization version (as also done in our follow-up work [18]).

(MAXIMUM-CARDINALITY) MATCHING

Input: An undirected graph $G = (V, E)$ and a nonnegative integer s .

Question: Is there a size- s subset $M_G \subseteq E$ of nonoverlapping (i.e. disjoint) edges?

Note that, for any polynomial-time solvable problem, solving the given instance and returning a trivial yes- or no-instance always produces a constant-size kernel in polynomial time. Hence, we are looking for kernelization algorithms that are *faster* than the algorithms solving the problem. The best we can hope for is linear time. For NP-hard problems, each polynomial-time kernelization algorithm is faster than any solution algorithm, unless $P=NP$. While the focus of classical kernelization for NP-hard problems is mostly on improving the *size* of the kernel, we particularly emphasize that for polynomially solvable problems it is mandatory to also focus on the *running time* of the kernelization algorithm. Indeed, we can consider linear-time kernelization as the holy grail and this drives our research when studying kernelization for MATCHING.

Our contributions. We present three kernels for MATCHING (see Table 1 for an overview). All our parameterizations can be categorized as “distance to triviality”

[5, 13]. They are motivated as follows. First, note that it is important that the parameters we exploit can be computed, or well approximated (within constant factors), in linear time regardless of the parameter value. Next, note that maximum-cardinality matchings can be trivially found in linear time on trees (or forests). That is why we consider the edge deletion distance (feedback edge number) and vertex deletion distance (feedback vertex number) to forests. Notably, there is a trivial linear-time algorithm for computing the feedback edge number and there is a linear-time factor-4 approximation algorithm for the feedback vertex number [1]. We mention in passing that the parameter vertex cover number, which is lower-bounded by the feedback vertex number, has been frequently studied for kernelization. In particular, Gianopoulou, Mertzios, and Niedermeier [12, 14] provided a linear-time computable quadratic-size kernel for MATCHING with respect to the parameter solution size (or equivalently vertex cover number). Coming to bipartite graphs, we parameterize by the vertex deletion distance to chain graphs which is motivated as follows. First, chain graphs form one of the most obvious easy cases for bipartite graphs where MATCHING can be solved in linear time [25]. Second, we show that the vertex deletion distance of any bipartite graph to a chain graph can be 4-approximated in linear time. Moreover, vertex deletion distance to chain graphs lower-bounds the vertex cover number of a bipartite graph.

An overview of our main results is given in Table 1. We study kernelization for MATCHING parameterized by the feedback vertex number, that is, the vertex deletion distance to a forest (see Sect. 3). As a warm-up we first show that a subset of our data reduction rules for the “feedback vertex set kernel” also yields a linear-time computable linear-size kernel for the typically much larger parameter feedback edge number (see Sect. 3.1). As for BIPARTITE MATCHING no faster algorithm is known than on general graphs, we kernelize BIPARTITE MATCHING with respect to the vertex deletion distance to chain graphs (see Sect. 4).

Seen from a high level, our two main results (Theorems 2 and 3, see Table 1) employ the same algorithmic strategy, namely upper-bounding (as a function of the parameter) the number of neighbors in the appropriate vertex deletion set (also called modulator) X ; that is, in the feedback vertex set or in the deletion set to chain graphs, respectively. To achieve this we develop new “irrelevant edge techniques” tailored to these two kernelization problems. More specifically, whenever a vertex v of the deletion set X has large degree, then we efficiently detect edges incident to v whose removal does not change the size of the maximum matching. Then the remaining graph can be further shrunk by scenario-specific data reduction rules. While this approach of removing irrelevant edges is natural, the technical details and the proofs of correctness become quite technical and combinatorially challenging. Note that there exists a trivial $O(km)$ -time solving (not only kernelization) algorithm, where k is the feedback vertex number. Our kernel has size $2^{O(k)}$. Therefore, only if $k = o(\log n)$ our kernelization algorithm provably shrinks the initial instance. However, our result is still relevant: First, our data reduction rules might assist in proving a polynomial upper bound—so our result is a first step in this direction. Second, the running time $O(kn)$ of our kernelization algorithm is a kind of “half way” between $O(km)$ (which could be as bad as $O(k^2n)$) and $O(n + m)$ (which is best possible). Finally, note that this work focuses on theoretical and worst-case analysis;

in practice, our kernelization algorithm might achieve much better upper bounds on real-world input instances. In fact, in experiments using the kernelization with respect to the feedback edge number, the observed kernels were always significantly smaller than the theoretical bound [18].

As a technical side remark, we emphasize that in order to achieve a linear-time kernelization algorithm, we often need to use suitable data structures and to carefully design the appropriate data reduction rules to be exhaustively applicable in linear time, making this form of “algorithm engineering” much more relevant than in the classical setting of mere polynomial-time data reduction rules.

2 Preliminaries and Basic Observations

Notation and Observations. We use standard notation from graph theory. A *feedback vertex (edge) set* of a graph G is a set X of vertices (edges) such that $G - X$ is a tree or forest. The *feedback vertex (edge) number* denotes the size of a minimum feedback vertex (edge) set. All paths we consider are simple paths. Two paths in a graph are called *internally vertex-disjoint* if they are either completely vertex-disjoint or they overlap only in their endpoints. A *matching* in a graph is a set of pairwise disjoint edges. Let $G = (V, E)$ be a graph and let $M \subseteq E$ be a matching in G . The degree of a vertex is denoted by $\deg(v)$. A vertex $v \in V$ is called *matched* with respect to M if there is an edge in M containing v , otherwise v is called *free* with respect to M . If the matching M is clear from the context, then we omit “with respect to M ”. An *alternating path* with respect to M is a path in G such that every second edge of the path is in M . An *augmenting path* is an alternating path whose endpoints are free. It is well known that a matching M is maximum if and only if there is no augmenting path for it. Let $M \subseteq E$ and $M' \subseteq E$ be two matchings in G . We denote by $G(M, M') := (V, M \triangle M')$ the graph containing only the edges in the symmetric difference of M and M' , that is, $M \triangle M' := M \cup M' \setminus (M \cap M')$. Observe that every vertex in $G(M, M')$ has degree at most two.

For a matching $M \subseteq E$ for G we denote by $M_G^{\max}(M)$ a maximum matching in G with the largest possible overlap (in number of edges) with M . That is, $M_G^{\max}(M)$ is a maximum matching in G such that for each maximum matching M' for G it holds that $|M \triangle M'| \geq |M \triangle M_G^{\max}(M)|$. Observe that if M is a maximum matching for G , then $M_G^{\max}(M) = M$. Furthermore observe that $G(M, M_G^{\max}(M))$ consists of only odd-length paths and isolated vertices, and each of these paths is an augmenting path for M . Moreover the paths in $G(M, M_G^{\max}(M))$ are as short as possible:

Observation 1 For any path v_1, v_2, \dots, v_p in $G(M, M_G^{\max}(M))$ it holds that $\{v_{2i-1}, v_{2j}\} \notin E$ for every $1 \leq i < j \leq p/2$.

Proof Assume that $\{v_{2i-1}, v_{2j}\} \in E$. Then $v_1, v_2, \dots, v_{2i-2}, v_{2i-1}, v_{2j}, v_{2j+1}, \dots, v_p$ is a shorter path which is also an augmenting path for M in G . The corresponding maximum matching M' satisfies $|M \triangle M_G^{\max}(M)| > |M \triangle M'|$, a contradiction to the definition of $M_G^{\max}(M)$. □

It is easy to see that removing k vertices in a graph can reduce the maximum matching size by at most k :

Observation 2 Let $G = (V, E)$ be a graph with a maximum matching M_G , let $X \subseteq V$ be a vertex subset of size k , and let M_{G-X} be a maximum matching for $G - X$. Then, $|M_{G-X}| \leq |M_G| \leq |M_{G-X}| + k$.

Kernelization. A *parameterized problem* is a set of instances (I, k) where $I \in \Sigma^*$ for a finite alphabet Σ , and $k \in \mathbb{N}$ is the *parameter*. We say that two instances (I, k) and (I', k') of parameterized problems P and P' are *equivalent* if (I, k) is a yes-instance for P if and only if (I', k') is a yes-instance for P' . A *kernelization* is an algorithm that, given an instance (I, k) of a parameterized problem P , computes in polynomial time an equivalent instance (I', k') of P (the *kernel*) such that $|I'| + k' \leq f(k)$ for some computable function f . We say that f measures the *size* of the kernel, and if $f(k) \in k^{O(1)}$, then we say that P admits a polynomial kernel. Typically, a kernel is achieved by applying polynomial-time executable data reduction rules. We call a data reduction rule \mathcal{R} *correct* if the new instance (I', k') that results from applying \mathcal{R} to (I, k) is equivalent to (I, k) . An instance is called *reduced* with respect to some data reduction rule if further application of this rule has no effect on the instance.

3 Kernelization for Matching on General Graphs

In this section, we investigate the possibility of efficient and effective preprocessing for MATCHING. As a warm-up, we first present in Sect. 3.1 a simple, linear-size kernel for MATCHING with respect to the parameter feedback edge number. Exploiting the data reduction rules and ideas used for this kernel, we then present in Sect. 3.2 the main result of this section: an exponential-size kernel for the almost always significantly smaller parameter feedback vertex number.

3.1 Warm-Up: Parameter Feedback Edge Number

We provide a linear-time computable linear-size kernel for MATCHING parameterized by the feedback edge number, that is, the size of a minimum feedback edge set. Observe that a minimum feedback edge set can be computed in linear time via a simple depth-first search or breadth-first search. The kernel is based on the next two simple data reduction rules due to Karp and Sipser [17]. These rules deal with vertices of degree at most two.

Reduction Rule 1 Let $v \in V$. If $\deg(v) = 0$, then delete v . If $\deg(v) = 1$, then delete v and its neighbor and decrease the solution size s by one (v is matched with its neighbor).

Reduction Rule 2 Let v be a vertex of degree two and let u, w be its neighbors. Then remove v , merge u and w , and decrease the solution size s by one.

The correctness was stated by Karp and Sipser [17]. For the sake of completeness, we give a proof.

Lemma 1 *Reduction Rules 1 and 2 are correct.*

Proof If v has degree zero, then clearly v cannot be in any matching and we can remove v .

If v has degree one, then let u be its single neighbor. Let M be a maximum matching of size at least s for G . Then u is matched in M since otherwise adding the edge $\{u, v\}$ would increase the size of the matching. Thus, a maximum matching in $G' = G - u - v$ has size at least $s - 1$. Conversely, a maximum matching of size $s - 1$ in G' can easily be extended by the edge $\{u, v\}$ to a maximum matching of size s in G .

If v has degree two, then let u and w be its two neighbors. Let M be a maximum matching of size at least s . If v is matched in M (i.e. either with the edge $\{u, v\}$ or with the edge $\{v, w\}$), then deleting v and merging u with w decreases the size of M by one. Similarly, if v is not matched in M , then both u and w are matched in M , since otherwise adding the edge $\{u, v\}$ (resp. $\{v, w\}$) would increase the size of the matching, a contradiction. Thus, in this case, deleting v and merging u with w decreases again the size of M by one (M loses either the edge incident to v or one of the edges incident to u and w). Hence, the resulting graph G'' has a maximum matching of size at least $s - 1$. Conversely, let M'' be a matching of size at least $s - 1$ for G'' . If the merged vertex vw is free, then $M := M'' \cup \{\{u, v\}\}$ is a matching of size s in G . Otherwise, vw is matched to some vertex y in M'' . Then matching y in G with either v or w (at least one of the two vertices is a neighbor of y) and matching u with the other vertex yields a matching of size at least s for G . □

While it is easy to exhaustively apply Reduction Rule 1 in linear time, applying Reduction Rule 2 exhaustively in linear time is nontrivial [2]. Note that applying Reduction Rule 2 might create new degree-one vertices and thus Reduction Rule 1 might become applicable again. To show our problem kernel, in the following theorem it is sufficient to first apply Reduction Rule 1 exhaustively and afterwards apply Reduction Rule 2 exhaustively.

Theorem 1 *MATCHING admits a linear-time computable linear-size kernel with respect to the parameter feedback edge number k .*

Proof Let G be the input graph. First we apply Reduction Rule 1 to G exhaustively, obtaining graph $G_1 = (V_1, E_1)$, and then we apply Reduction Rule 2 to G_1 exhaustively, obtaining graph $G_2 = (V_2, E_2)$. Note that both G_1 and G_2 can be computed in linear time [2]. We will prove that G_2 has at most $6k$ vertices and $7k$ edges. Denote with $X_1 \subseteq E_1$ and $X_2 \subseteq E_2$ the minimum feedback edge sets for G_1 and G_2 respectively. Note that $|X_2| \leq |X_1| \leq k$. For any graph H , denote with V_H^1 , V_H^2 , and $V_H^{\geq 3}$ the vertices of H that have degree one, two, and more than two, respectively (in our case H will be replaced by $G_1 - X_1$ or $G_2 - X_2$, respectively). Observe that all vertices

in G_1 have degree at least two, since G_1 is reduced with respect to Reduction Rule 1. Thus $|V_{G_1-X_1}^1| \leq 2k$, as each leaf in $G_1 - X_1$ has to be incident to an edge in X_1 . Next, since $G_1 - X_1$ is a forest, we have that $|V_{G_1-X_1}^{\geq 3}| < |V_{G_1-X_1}^1|$, and thus $|V_{G_1-X_1}^{\geq 3}| < 2k$. Note that the number of degree-two vertices in G_1 cannot be upper-bounded by a function of k . However, observe that the exhaustive application of Reduction Rule 2 to G_1 removes all vertices that have degree-two in G_1 and possibly merges some of the remaining vertices. Thus, G_2 contains no vertices with degree two and thus, $|V_{G_2-X_2}^2| \leq 2k$. Altogether, the number of vertices in G_2 is at most $|V_{G_1-X_1}^1| + |V_{G_2-X_2}^2| + |V_{G_1-X_1}^{\geq 3}| \leq 6k$. Since $G_2 - X_2$ is a forest, it follows that G_2 has at most $|V_2| + k \leq 7k$ edges. \square

Applying the $O(m\sqrt{n})$ -time algorithm for MATCHING [21] on the above kernel yields the following.

Corollary 1 *MATCHING can be solved in $O(n + m + k^{1.5})$ time, where k is the feedback edge number.*

3.2 Parameter Feedback Vertex Number

We next provide for MATCHING a kernel of size $2^{O(k)}$ computable in $O(kn)$ time where k is the feedback vertex number. Using a known linear-time factor 4-approximation algorithm [1], we can compute an approximate feedback vertex set and use it in our kernelization algorithm.

Roughly speaking, our kernelization algorithm extends the linear-time computable kernel with respect to the parameter feedback edge set. Thus, Reduction Rule 1 and 2 play an important role in the kernelization. Compared to the other kernels presented in this paper, the kernel presented here comes at the price of higher running time $O(kn)$ and bigger kernel size (exponential). It remains open whether MATCHING parameterized by the feedback vertex number admits a linear-time computable kernel (possibly of exponential size), or whether it admits a polynomial kernel computable in $O(kn)$ time.

Subsequently, we describe our kernelization algorithm which keeps in the kernel all vertices in the given feedback vertex set X and shrinks the size of $G - X$. Before doing so, we need some further notation. In this section, we assume that each tree is rooted at some arbitrary (but fixed) vertex such that we can refer to the parent and children of a vertex. A leaf in $G - X$ is called a *bottommost leaf* either if it has no siblings or if all its siblings are also leaves. (Here, bottommost refers to the subtree with the root being the parent of the considered leaf.) The outline of the algorithm is as follows (we assume throughout this section that $k < \log n$ since otherwise the input instance is already a kernel of size $O(2^k)$):

1. Reduce G wrt. Reduction Rule 2 and 1.

2. Compute a maximum matching M_{G-X} in $G - X$ (where X is a feedback vertex set that is computed by the linear-time 4-approximation algorithm [1]).
3. Modify M_{G-X} in linear time such that only the leaves of $G - X$ are free.
4. Upper-bound the number of free leaves in $G - X$ by k^2 (Sect. 3.2.1).
5. Upper-bound the number of bottommost leaves in $G - X$ by $O(k^2 2^k)$ (Sect. 3.2.2).
6. Upper-bound the degree of each vertex in X by $O(k^2 2^k)$. Then, use Reduction Rule 2 and 1 to provide the kernel of size $2^{O(k)}$ (Sect. 3.2.3).

Whenever we reduce the graph at some step, we also show that the reduction is correct. That is, the given instance is a yes-instance if and only if the reduced one is a yes-instance. The correctness of our kernelization algorithm then follows by the correctness of each step. We discuss in the following the details of each step.

3.2.1 Steps 1–4

In this subsection, we first discuss the straightforward Steps 1 to 3 and then turn to Step 4.

Steps 1–3. As in Sect. 3.1, we perform Step 1 in linear time by first applying Reduction Rule 1 and then Reduction Rule 2 using the algorithm due to Bartha and Kresz [2]. By Lemma 1 this step is correct.

A maximum matching M_{G-X} in Step 2 can be computed by repeatedly matching a free leaf to its neighbor and by removing both vertices from the graph (thus effectively applying Reduction Rule 1 to $G - X$). Clearly, this can be done in linear time.

Step 3 can be done in $O(n)$ time by traversing each tree in $G - X$ in a BFS manner starting from the root: If a visited inner vertex v is free, then observe that all children are matched since M_{G-X} is maximum. Pick an arbitrary child u of v and match it with v . The vertex w that was previously matched to u is now free and since it is a child of u , it will be visited in the future. Observe that Steps 2 and 3 do not change the graph but only the auxiliary matching M_{G-X} , and thus the first three steps are correct. The next observation summarizes the short discussion above.

Observation 3 Steps 1 to 3 are correct and can be applied in linear time.

Step 4. Our goal is to upper-bound the number of edges between vertices of X and $V \setminus X$, since we can then use a simple analysis as for the parameter feedback edge set. Furthermore, recall that by Observation 2 the size of any maximum matching in G is at most k plus the size of M_{G-X} . Now, the crucial observation here is that, if a vertex $x \in X$ has at least k neighbors $\{v_1, \dots, v_k\}$ in $V \setminus X$ which are free wrt. M_{G-X} , then there exists a maximum matching where x is matched to one of $\{v_1, \dots, v_k\}$ since at most $k - 1$ can be “blocked” by other matching edges. Indeed, consider otherwise a maximum matching M in which x is not matched with any of $\{v_1, \dots, v_k\}$. Then, since $|X| = k$, note that at most $k - 1$ vertices among $\{v_1, \dots, v_k\}$ are matched in M with a vertex in X ; suppose without loss of generality that v_k is not matched with any vertex in X (and thus v_k is not matched at all in M). If x is unmatched in M , then the matching $M \cup \{x, v_k\}$ has greater cardinality than M , a contradiction. Otherwise, if

x is matched in M with a vertex z , then $M \cup \{\{x, v_k\}\} \setminus \{\{x, z\}\}$ is another maximum matching of G , in which x is matched with a vertex among $\{v_1, \dots, v_k\}$. Formalizing this idea, we obtain the following data reduction rule.

Reduction Rule 3 Let $G = (V, E)$ be a graph, let $X \subseteq V$ be a vertex subset of size k , and let M_{G-X} be a maximum matching for $G - X$. If there is a vertex $x \in X$ with at least k free neighbors $V_x = \{v_1, \dots, v_k\} \subseteq V \setminus X$, then delete all edges from x to vertices in $V \setminus V_x$.

We first show the correctness and then the running time of Reduction Rule 3.

Lemma 2 *Reduction Rule 3 is correct.*

Proof Denote by s the size of a maximum matching in the input graph $G = (V, E)$ and by s' the size of a maximum matching in the new graph $G' = (V', E')$, where some edges incident to x are deleted. We need to show that $s = s'$. Since any matching in G' is also a matching in G , we easily obtain $s \geq s'$.

It remains to show $s \leq s'$. To this end, let $M_G := M_G^{\max}(M_{G-X})$ be a maximum matching for G with the maximum overlap with M_{G-X} (see Sect. 2). If x is free wrt. M_G or if x is matched to a vertex v that is also in G' a neighbor of x , then M_G is also a matching in G' ($M_G \subseteq E'$) and thus we have $s \leq s'$. Hence, consider the remaining case where x is matched to some vertex v such that $\{v, x\} \notin E'$, that is, the edge $\{v, x\}$ was deleted by Reduction Rule 3. Hence, x has k neighbors v_1, \dots, v_k in $V \setminus X$ such that each of these neighbors is free wrt. M_{G-X} and none of the edges $\{v_i, x\}, i \in [k]$, was deleted. Observe that by the choice of M_G , the graph $G(M_{G-X}, M_G)$ (the graph over vertex set V and the edges that are either in M_{G-X} or in M_G , see Sect. 2) contains exactly $s - |M_{G-X}|$ paths of length at least one. Each of these paths is an augmenting path for M_{G-X} . By Observation 2, we have $s - |M_{G-X}| \leq k$. Observe that $\{v, x\}$ is an edge in one of these augmenting paths; denote this path with P . Thus, there are at most $k - 1$ augmenting paths in $G(M_{G-X}, M_G)$ that do not contain x . Also, each of these paths contains exactly two vertices that are free wrt. M_{G-X} : the endpoints of the path. This means that no vertex in X is an inner vertex on such a path. Furthermore, since M_{G-X} is a maximum matching, it follows that for each path at most one of these two endpoints is in $V \setminus X$. Hence, at most $k - 1$ vertices of v_1, \dots, v_k are contained in the $k - 1$ augmenting paths of $G(M_{G-X}, M_G)$ except P . Consequently, one of these vertices, say v_i , is free wrt. M_G and can be matched with x . Thus, by reversing the augmentation along P and adding the edge $\{v_i, x\}$ we obtain another matching M'_G of size s . Observe that M'_G is a matching for G and for G' and thus we have $s \leq s'$. This completes the proof of correctness. \square

Lemma 3 *Reduction Rule 3 can be exhaustively applied in $O(n + m)$ time.*

Proof We exhaustively apply the data reduction rule as follows. First, initialize for each vertex $x \in X$ a counter with zero. Second, iterate over all free vertices in $G - X$ in an arbitrary order. For each free vertex $v \in V \setminus X$ iterate over its neighbors in X . For each neighbor $x \in X$ do the following: if the counter is less than k , then increase the counter by one and mark the edge $\{v, x\}$ (initially all edges are unmarked). Third, iterate over all vertices in X . If the counter of the currently considered vertex x is k , then delete all unmarked edges incident to x . This completes the algorithm. Clearly, it deletes edges incident to a vertex $x \in X$ if and only if x has k free neighbors in $V \setminus X$ and the edges to these k neighbors are kept. The running time is $O(n + m)$: When iterating over all free vertices in $V \setminus X$ we consider each edge at most once. Furthermore, when iterating over the vertices in X , we again consider each edge at most once. □

To finish Step 4, we exhaustively apply Reduction Rule 3 in linear time. Afterwards, there are at most k^2 free (wrt. to M_{G-X}) leaves in $G - X$ that have at least one neighbor in X since each of the k vertices in X is adjacent to at most k free leaves. Thus, applying Reduction Rule 1 we can remove the remaining free leaves that have no neighbor in X . However, since for each degree-one vertex also its neighbor is removed, we might create new free leaves in $G - X$ and need to again apply Reduction Rule 3 and update the matching (see Step 3). This process of alternating application of Reduction Rule 1 and 3 stops after at most k rounds since the neighborhood of each vertex in X can be changed by Reduction Rule 3 at most once. This shows the running time $O(k(n + m))$. We next show how to improve this to $O(n + m)$. In doing so, we arrive at the central proposition of this subsection, stating that Steps 1 to 4 can be performed in linear time.

Proposition 1 *Given a matching instance (G, s) and a feedback vertex set X , Algorithm 1 computes in linear time an instance (G', s') with feedback vertex set X and a maximum matching $M_{G'-X}$ in $G' - X$ such that the following holds.*

- *There is a matching of size s in G if and only if there is a matching of size s' in G' .*
- *Each vertex in $G' - X$ that is free wrt. $M_{G'-X}$ is a leaf in $G' - X$.*
- *There are at most k^2 free leaves in $G' - X$.*

Algorithm 1: An algorithm performing Steps 1 to 4 in linear time.

```

Input: A matching instance  $(G = (V, E), s)$  and a feedback vertex set  $X \subseteq V$  for  $G$ 
with  $|X| = k$ .
Output: An equivalent matching instance  $(G', s')$  such that  $X$  is also a feedback
vertex set for  $G'$  and a maximum matching  $M_{G'-X}$  for  $G' - X$  such that
only at most  $k^2$  leaves in  $G' - X$  are free with respect to  $M_{G'-X}$ .

1 Reduce  $G$  wrt. Reduction Rules 1 and 2
2 Compute a maximum matching  $M_{G-X}$  as described in Step 3 // see Observation 3
3 foreach  $x \in X$  do  $c(x) \leftarrow 0$  //  $c(x)$  will store the number of free neighbors for  $x$ 
4 foreach  $e \in E$  do  $\text{marked}(e) \leftarrow \text{False}$ 
5  $L \leftarrow$  stack containing all free leaves in  $G - X$  in any order
6 while  $L$  is not empty do
7    $u \leftarrow \text{pop}(L)$ 
8   foreach  $x \in N_G(u) \cap X$  do // Check if Reduction Rule 3 is applicable for  $x$ 
9      $c(x) \leftarrow c(x) + 1$ ,  $\text{marked}(\{u, x\}) \leftarrow \text{True}$  // fix  $u$  as free neighbor of  $x$ 
10    if  $c(x) = k$  then //  $x$  has enough free neighbors: apply Reduction Rule 3
11      foreach  $y \in N_G(x) \cap X$  do delete  $\{x, y\}$ 
12      foreach  $v \in N_G(x) \setminus X$  do
13        if  $\text{marked}(\{x, v\}) = \text{False}$  then
14          delete  $\{x, v\}$ 
15          /* Next deal with the case that  $v$  is a free leaf in  $G - X$  and  $x$ 
was the last neighbor of  $v$  in  $X$  */
16          if  $\text{deg}_G(v) = \text{deg}_{G-X}(v) = 1$  and  $v$  is free then push  $v$  on  $L$ 
17 if  $\text{deg}_G(u) = \text{deg}_{G-X}(u) = 1$  then //  $u$  has no neighbors in  $x$ 
18    $v \leftarrow$  neighbor of  $u$  in  $G - X$ ;  $w \leftarrow$  matched neighbor of  $v$ 
19   delete  $u$  and  $v$  from  $G$  // apply Reduction Rule 1
20    $M_{G-X} \leftarrow M_{G-X} \setminus \{\{w, v\}\}$ ,  $s \leftarrow s - 1$  // update  $M_{G-X}$  and  $s$ 
21   /* augment along an arbitrary alternating path from  $w$  to a leaf in the
subtree rooted in  $w$ : */
22   while  $w$  is not a leaf in  $G - X$  do
23      $w' \leftarrow$  arbitrary child of  $w$ ;  $w'' \leftarrow$  matched neighbor of  $w'$ 
24      $M_{G-X} \leftarrow (M_{G-X} \setminus \{\{w', w''\}\}) \cup \{\{w, w'\}\}$ 
25      $w \leftarrow w''$ 
26   push  $w$  to  $L$  //  $w$  is a free leaf, so add  $w$  to the list of vertices to check
27 return  $(G, s)$  and  $M_{G-X}$ .

```

Before proving Proposition 1, we explain Algorithm 1 which reduces the graph with respect to Reduction Rules 1 and 3 and updates the matching M_{G-X} as described in Step 3. The algorithm performs in Lines 1 and 2 Steps 1 to 3. This can be done in linear time (see Observation 3). Next, Reduction Rule 3 is applied in Lines 8 to 15 using the approach described in the proof of Lemma 3: For each vertex in x a counter $c(x)$ is maintained. When iterating over the free leaves in $G - X$, these counters will be updated. If a counter $c(x)$ reaches k , then the algorithm knows that x has k fixed free neighbors and according to Reduction Rule 3 the edges to all other vertices can be deleted (see Line 10). Observe that once the counter $c(x)$ reaches k , the vertex x will never be considered again by the algorithm since its only remaining neighbors are free leaves in $G - X$ that




Fig. 1 Dealing with new degree-one vertices occurring during the application of Reduction Rule 3 within Step 4. Only vertices visited in the tree $G - X$ in Lines 16–24 of Algorithm 1 are shown. Further possible neighbors are indicated by edges. Left side: Vertex v is a free leaf in $G - X$ (vertices in X are not illustrated). The gray highlighted alternating path indicates where Algorithm 1 augments the maximum matching M_{G-X} in $G - X$. Bold edges indicate edges in M_{G-X} . Vertex w'' is the leaf where the augmentation stops (w'' is matched, otherwise M_{G-X} would not be a maximum matching). Right side: Situation after Algorithm 1 augmentation. Vertex w'' will be added to the list L and further processed

already have been popped from the stack L . The only difference from the description in the proof of Lemma 3 is that the algorithm reacts if some leaf v in $G - X$ lost its last neighbor in X (see Line 15). If v is free, then add v to the stack L of unmatched degree-one vertices and defer dealing with v to a second stage of the algorithm (in Line 16 to 24). (If v is matched, then we deal with v in Step 6.)

We next discuss this second stage from Lines 16–24 (see Fig. 1 for an illustration): Let u be an entry in L such that u has degree one in Line 16, that is, u is a free leaf in $G - X$ and has no neighbors in X . Then, following Reduction Rule 1, delete u and its neighbor v and decrease the solution size s by one (see Lines 18 and 19). Let w denote the previously matched neighbor of v . Since v was removed, w is now free. If w is a leaf in $G - X$, then we can simply add it to L and deal with it later. If w is not a leaf, then we need to update M_{G-X} since only leaves are allowed to be free. To this end, augment along an arbitrary alternating path from w to a leaf in the subtree with root w (see Lines 20 to 23). This is done as follows: Pick an arbitrary child w' of w . Let w'' be the matched neighbor of w' . Observe that w'' has to exist as if w' would be free, then $\{w, w'\}$ could be added to M_{G-X} ; a contradiction to the maximality of M_{G-X} . Since w is the parent of w' , it follows that w'' is a child of w' . Now, remove $\{w', w''\}$ from M_{G-X} , add $\{w', w\}$ and repeat the procedure with w'' taking the role of w . Observe that the endpoint of this found alternating path, after augmentation, always is a free leaf. Thus, this free leaf needs to be pushed to L . This completes the algorithm description.

The correctness of Algorithm 1 (stated in the next lemma) follows in a straightforward way from the above discussion. For the formal proofs we introduce some notation. We denote by G_i (respectively M_i) the intermediate graph (respectively matching) stored by Algorithm 1 before the i^{th} iteration of the while loop in Line 6,

that is, G_1 is the input graph and M_1 is the initial matching computed in Line 2. The following observation is easy to see but useful in our proofs.

Observation 4 For each $i \in \{1, \dots, q\}$ where q is the number of iterations of the while loop in Line 6, we have that M_i is a maximum matching for $G_i - X$. If $i \geq 2$, then G_i is a subgraph of G_{i-1} .

Lemma 4 Algorithm 1 is correct, that is, given a matching instance (G, s) and a feedback vertex set X , it computes an instance (G', s') with feedback vertex set X and a maximum matching $M_{G'-X}$ in $G' - X$ such that:

1. There is a matching of size s in G if and only if there is a matching of size s' in G' .
2. Each vertex in $G' - X$ that is free wrt. $M_{G'-X}$ is a leaf in $G' - X$.
3. There are at most k^2 free vertices in $G' - X$.

Proof Observation 4 implies that the returned graph G' is a subgraph of the input graph G . Thus, X is a feedback vertex set for both these graphs. Moreover, by Observation 4, $M_{G'-X}$ is a maximum matching for $G' - X$.

As to 1, observe that Algorithm 1 obtains G' from G by deleting edges in Line 14 according to Reduction Rule 3 and by deleting vertices in Line 18 according to Reduction Rule 1. Thus, 1 follows from the correctness of these data reduction rules (see Lemmas 1 and 2).

As to 2, observe that $G - X$ is changed if and only if the matching M_{G-X} is changed accordingly (see Lines 16 to 24). That is, after each deletion of vertices, the algorithm ensures that only leaves are free. Moreover, during the algorithm M_{G-X} is always a maximum matching for $G - X$.

As to 3, observe that any free leaf in $G - X$ that is not removed needs to have a neighbor in X (see Line 16). As Reduction Rule 3 is applied in Lines 8 to 15, there are at most k^2 such free leaves. \square

We next show that Algorithm 1 runs in linear time. To this end, we need a further technical statement.

Lemma 5 In G_i , let P be an even-length alternating path wrt. M_i from a free leaf r to a matched inner vertex t of $G_i - X$. Let u be the matched neighbor of t . Then for each $j \in \{1, \dots, i\}$ there exists in G_j an even-length alternating path P' from t to a free leaf r' such that the neighbor of t on P' is either (i) u , (ii) t 's parent, or (iii) a vertex not contained in G_i .

Proof We prove the statement of the lemma by induction on i . The base case $i = 1$ is trivial since $G_1 = G$ and thus $P' = P$.

Now assume the statement is true for G_{i-1} , $i \geq 2$. We show that it holds for G_i as well. By Observation 4, G_i is a subgraph of G_{i-1} (and of G). Thus, the path P is also contained in G_{i-1} (and in G). If r is a leaf in $G_{i-1} - X$ and if M_i contains the same edges of P as M_{i-1} , then P is an even-length augmenting path in G_{i-1} and the

statement of the lemma follows from applying the induction hypothesis and Observation 4. Thus, assume that (a) r is not a leaf in $G_{i-1} - X$ or (b) M_i does not contain the same edges of P as M_{i-1} (or both).

We start with case (a) assuming that r is not a leaf in $G_{i-1} - X$. Then in the $(i - 1)$ st iteration of the while loop in Line 6, Algorithm 1 deleted the child r' of r and the child r'' of r' in Line 18. Moreover, M_{i-1} contained the edge $\{r, r'\}$ and r'' was a free leaf in $G_{i-1} - X$. Thus, extending P by the two vertices r', r'' yields in G_{i-1} an even-length alternating path P^* from t to the free leaf r'' such that the neighbor of t on P^* is u . Hence, the statement of the lemma follows from the induction hypothesis and Observation 4.

We next consider case (b), assuming that M_i and M_{i-1} do not contain the same edges of P . Thus, in the $(i - 1)$ st iteration of the while loop in Line 6, Algorithm 1 augmented along some alternating path in Lines 20 to 23. Denote with Q this alternating path and let w_q be starting point of Q , that is, w_q is the vertex w in Line 17. Let v_Q, u_Q be the two deleted vertices in Line 17. Let r_Q be the other endpoint of Q , that is, r_Q is a leaf in G_{i-1} and thus a free leaf in G_i . Since M_i and M_{i-1} differ on P , this implies that the two paths Q and P overlap. Let z be the vertex on P and on Q which is closest to r . If $z = r = r_q$, then P is a subpath of Q and in G_{i-1} there is an alternating path P^* from t to the free leaf u_Q . (Here, P^* is the part of Q that is not contained in P .) Since the alternating path built in Line 20 to 23 is only extended by selecting child vertices, this implies that $w_q = t$ or w_q is an ancestor of t . Thus, the neighbor of t in P^* is either t 's parent or v_Q , that is, a child of t not contained in G_i . Hence, the statement of the lemma follows from the induction hypothesis and Observation 4.

It remains to consider the case that $z \neq r$. Let $z_Q (z_p)$ be the neighbor of z that is on Q but not on P (on P but not on Q); similarly let z_{pQ} be the neighbor of z that is on both P and Q . Since Q is an alternating path either $\{z, z_Q\}$ or $\{z, z_{pQ}\}$ is in M_{i-1} .

First consider the case that $\{z, z_Q\}$ is in M_{i-1} . Then, since both the subpath of Q from z to u_Q and the subpath of P from z to r are alternating, we obtain an augmenting path from u_Q over z to r . This is a contradiction to the maximality of M_{i-1} .

Second, consider the case that $\{z, z_{pQ}\}$ is in M_{i-1} . Thus, (after augmenting Q) the edge $\{z, z_{pQ}\}$ is not in M_i . Moreover, as $\{z, z_Q\}$ is in M_i , the edge $\{z, z_p\}$ is also not in M_i . This contradicts the fact that P is an alternating path. □

Lemma 6 *Algorithm 1 runs in $O(n + m)$ time.*

Proof By Observation 3, Steps 1 to 3 in Lines 1 and 2 can be executed in linear time. Moreover, it is easy to execute Lines 3 to 5 in one sweep over the graph, that is, in linear time. It remains to show that Lines 6 to 24 run in linear time. To this end, we prove that each edge in E is being processed at most two times in Lines 6 to 24.

Start with the edges with at least one endpoint in X . These edges will be inspected at most twice by the algorithm: Once, when the edge is marked (see Line 9). The second time is when the edge is checked and possibly deleted Lines 13 and 14. This shows that the first part (Lines 8 to 15) runs in linear time.

It remains to consider the edges within $G - X$. To this end, observe that the algorithm performs two actions on the edges: deleting the edges (Line 18) and finding and augmenting along an alternating path (Lines 20 to 23). Clearly, after deleting an edge it will no longer be considered, so it remains to show that each edge is part of at most one alternating path in Line 22. Assume toward a contradiction that the algorithm augments along an edge twice or more. From all the edges that are augmented twice or more let $e \in E$ be one that is closest to the root of the tree containing e , that is, there is no edge closer to a root. Let P_1 and P_2 be the first two augmenting paths containing e . Assume without loss of generality that the algorithm augmented along P_1 in iteration i_1 and along P_2 in iteration i_2 of the while loop in Line 6 with $i_1 \leq i_2$. Let w_1 and w_2 be the two start points (the respective vertex w in Line 17) of P_1 and P_2 respectively. Let u_1 and v_1 (u_2 and v_2) be the vertices deleted in Line 18 which in turn made w_1 (w_2) free. Observe that e does not contain any of these four vertices u_1, v_1, u_2, v_2 since before augmenting P_1 (P_2) the vertices u_1 and v_1 (u_2 and v_2) are deleted in Line 18. Since e is contained in both paths, either w_1 is an ancestor of w_2 or vice versa (or $w_1 = w_2$).

Assume first that w_2 is an ancestor of w_1 . Thus, $e = \{w_1, w'_1\}$ where $w'_1 \neq v_1$ and w'_1 is a child of w_1 (see Line 21). Consider G_{i_2} and M_{i_2} before the augmentation along P_2 . Clearly, in G_{i_2} there is an alternating path of length two from w_2 to the free leaf u_2 . Thus, by Lemma 5, in G_{i_1} there is an alternating path Q_1 from w_2 to a free leaf r such that r and w_1 are not in the same subtree of w_2 . Moreover, by choice of e the two matchings M_{i_1} and M_{i_2} contain the same edges on the path from w_1 to w_2 in $G - X$. Hence, there is an alternating path Q_2 from w_1 to w_2 in G_{i_1} . There is also an alternating path Q_3 from w_1 to the free leaf u_1 in G_{i_1} (see Line 17). Combining Q_1, Q_2, Q_3 gives an augmenting path from u_1 to r in G_{i_1} ; a contradiction to the maximality of M_{i_1} (see Observation 4).

Next, consider the case that $w_1 = w_2$. By choice of e we have that $e = \{w_1, w'\}$ with w' being a child of w_1 in G_{i_2} and $w' \neq v_2$. Thus, after the augmentation along P_1 the edge e is matched (see Line 21). This is a contradiction to the choice of P_2 and the fact that $\{w_2, v_2\} \in M_{i_2}$ (see Line 17).

Finally, consider the case that w_1 is an ancestor of w_2 . By choice of e we have that $e = \{w_2, w'_2\}$ with w'_2 being a child of w_2 in G_{i_2} and $w'_2 \neq v_2$. From the argumentation used in the case $w_1 = w_2$ above, we can infer that after augmenting P_1 the edge e is not matched, thus $e \notin M_{i_1+1}$ and $e \in M_{i_1}$. Observe that in G_{i_2} there is a length-two alternating path from w_2 to the free leaf u_2 . Thus, by Lemma 5, there is an even-length alternating path P from w_2 to a free leaf in G_i . Moreover, the (matched) neighbor w'_2 of w_2 in P is either (i) v_2 , (ii) the parent of w_2 , or (iii) a vertex not in G_{i_2} . Since $e \in M_{i_1}$, it follows that w'_2 is the matched neighbor of w_2 on P . However, w'_2 is in G_{i_2} , is neither the parent of w_2 nor of v_2 , a contradiction. \square

Proposition 1 now follows from Lemmas 4 and 6.

3.2.2 Step 5

In this step we reduce the graph in $O(kn)$ time so that at most $k^2(2^k + 1)$ bottommost leaves will remain in the forest $G - X$. We will restrict ourselves to consider leaves that are matched with their parent vertex in M_{G-X} and that do not have a sibling. We call these bottommost leaves *interesting*. Any sibling of a bottommost leaf is by

definition also a leaf. Thus, at most one of these leaves (the bottommost leaf or one of its siblings) is matched with respect to M_{G-X} and all other leaves are free. Recall that in the previous step we upper-bounded the number of free leaves with respect to M_{G-X} by k^2 . Hence there are at most $2k^2$ bottommost leaves that are not interesting (each free leaf can be a bottommost leaf with a sibling matched to the parent).

Our general strategy for this step is to extend the idea behind Reduction Rule 3: We want to keep for each pair of vertices $x, y \in X$ at most k different internally vertex-disjoint augmenting paths from x to y . In this step, we only consider augmenting paths of the form x, u, v, y where v is a bottommost leaf and u is v 's parent in $G - X$. Assume that the parent u of v is adjacent to some vertex $x \in X$. Observe that in this case any augmenting path starting with the two vertices x and u has to continue to v and end in a neighbor of v . Thus, the edge $\{x, u\}$ can be only used in augmenting paths of length three. Furthermore, for different parent vertices $u \neq u'$ the length-three augmenting paths are clearly internally vertex-disjoint. If we do not need the edge $\{x, u\}$ because we kept k augmenting paths from x to each neighbor $y \in N(v) \cap X$ already, then we can delete $\{x, u\}$. Furthermore, if we deleted the last edge from u to X (or u had no neighbors in X in the beginning), then u is a degree-two vertex in G and can be removed by applying Reduction Rule 2. As the child v of u is a leaf in $G - X$, it follows that v has at most $k + 1$ neighbors in G . We show below (Lemma 7) that the application of Reduction Rule 2 to remove u takes $O(k)$ time. As we remove at most n vertices, at most $O(kn)$ time is spent on Reduction Rule 2 in this step.

We now show that, after a simple preprocessing, one application of Reduction Rule 2 in the algorithm above can indeed be performed in $O(k)$ time.

Lemma 7 *Let u be a leaf in the tree $G - X$, v be its parent, and let w be the parent of v . If v has degree two in G , then applying Reduction Rule 2 to v (deleting v , merging u and v , and setting $s := s - 1$) can be done in $O(k)$ time plus $O(kn)$ time for an initial preprocessing.*

Proof The preprocessing is to simply create a partial adjacency matrix for G with the vertices in X in one dimension and V in the other dimension. This adjacency matrix has size $O(kn)$ and can clearly be computed in $O(kn)$ time.

Now apply Reduction Rule 2 to v . Deleting v takes constant time. To merge u and w iterate over all neighbors of u . If a neighbor u' of u is already a neighbor of w , then decrease the degree of u' by one, otherwise add u' to the neighborhood of w . Then, relabel w to be the new merged vertex uw .

Since u is a leaf in $G - X$ and its only neighbor in $G - X$, namely v , is deleted, it follows that all remaining neighbors of u are in X . Thus, using the above adjacency matrix, one can check in constant time whether u' is a neighbor of w . Hence, the above algorithm runs in $O(\deg(u)) = O(k)$ time. □

The above ideas are used in Algorithm 2 which we use for this step (Step 5). The algorithm is explained in the proof of the following proposition stating the correctness and the running time of Algorithm 2.

Algorithm 2: An algorithm performing Step 5 in $O(kn)$ time.

Input: A matching instance $(G = (V, E), s)$, a feedback vertex set $X \subseteq V$ of size k for G with $k < \log n$, and a maximum matching M_{G-X} for $G - X$ with at most k^2 free vertices in $G - X$ that are all leaves.

Output: An equivalent matching instance (G', s') such that X is also a feedback vertex set for G' and $G' - X$ is a tree with at most $k^2(2^k + 1)$ bottommost leaves, and a maximum matching $M_{G'-X}$ for $G' - X$ with at most k^2 free vertices in $G' - X$ that are all leaves.

```

1 Fix an arbitrary bijection  $f: 2^X \rightarrow \{1, \dots, 2^k\}$ 
2 foreach  $v \in V \setminus X$  do
3   | Set  $f_X(v) \leftarrow f(N(v) \cap X)$  // The number  $f_X(v) < n$  can be read in constant
   |   time.
4 Initialize a table Tab of size  $k \cdot 2^k$  with  $\text{Tab}[x, f(Y)] \leftarrow 0$  for all  $x \in X, \emptyset \subsetneq Y \subseteq X$ 
5  $P \leftarrow$  list containing all parents of interesting bottommost leaves
6 while  $P$  is not empty do
7    $u \leftarrow \text{pop}(P)$ 
8    $v \leftarrow$  child vertex of  $u$  in  $G - X$ 
9   foreach  $x \in N(u) \cap X$  do
10    | if  $\text{Tab}[x, f_X(v)] < k$  then
11    |   |  $\text{Tab}[x, f_X(v)] \leftarrow \text{Tab}[x, f_X(v)] + 1$ 
12    |   else
13    |     | delete  $\{x, u\}$ 
14    | if  $u$  has now degree two in  $G$  then
15    |   | Apply Reduction Rule 2 to  $u$  // This decreases  $s$  by one.
16    |   |  $vw \leftarrow$  vertex resulting from merging  $v$  and the parent  $w$  of  $u$ 
17    |   | if  $vw$  is now an interesting bottommost leaf then
18    |   |   | add the parent of  $vw$  to  $P$ 
19 return  $(G, s)$  and  $M_{G-X}$ .
```

Proposition 2 Let $(G = (V, E), s)$ be a MATCHING instance, let $X \subseteq V$ be a feedback vertex set, and let M_{G-X} be a maximum matching for $G - X$ with at most k^2 free vertices in $G - X$ that are all leaves. Then, Algorithm 2 computes in $O(kn)$ time an instance (G', s') with feedback vertex set X and a maximum matching $M_{G'-X}$ in $G' - X$ such that the following holds.

- There is a matching of size s in G if and only if there is a matching of size s' in G' .
- There are at most $2k^2(2^k + 1)$ bottommost leaves in $G' - X$.
- There are at most k^2 free vertices in $G' - X$ and they are all leaves.

Proof We start with describing the basic idea of the algorithm. To this end, let $\{u, v\} \in E$ be an edge such that v is an interesting bottommost leaf, that is, v has no siblings and is matched to its parent u by M_{G-X} . Counting for each pair $x \in N(u) \cap X$ and $y \in N(v) \cap X$ one augmenting path in a simple worst-case analysis gives $O(k^2)$ time per edge, which is too slow for our purposes. Instead, we count for each pair consisting of a vertex $x \in N(u) \cap X$ and a set $Y = N(v) \cap X$ one augmenting path. In this way, we know that for each $y \in Y$ there is one augmenting

path from x to y without iterating through all $y \in Y$. This comes at the price of considering up to $k2^k$ such pairs. However, we will show that we can do the computations in $O(k)$ time per considered edge in $G - X$. The main reason for this improved running time is a simple preprocessing that allows for a bottommost vertex v to determine $N(v) \cap X$ in constant time.

The preprocessing is as follows (see Lines 1 to 3): First, fix an arbitrary bijection f between the set of all subsets of X to the numbers $\{1, 2, \dots, 2^k\}$. This can be done for example by representing a set $Y \subseteq X = \{x_1, \dots, x_k\}$ by a length- k binary string (a number) where the i^{th} position is 1 if and only if $x_i \in Y$. Given a set $Y \subseteq X$ such a number can be computed in $O(k)$ time in a straightforward way. Thus, Lines 1 to 3 can be performed in $O(kn)$ time. Furthermore, since we assume that $k < \log n$ (otherwise the input instance has already at most 2^k vertices), we have that $f(Y) < n$ for each $Y \subseteq X$. Thus, reading and comparing these numbers can be done in constant time. Furthermore, in Line 3 the algorithm precomputes for each vertex the number corresponding to its neighborhood in X .

After the preprocessing, the algorithm uses a table Tab where it counts an augmenting path from a vertex $x \in X$ to a set $Y \subseteq X$ whenever a bottommost leaf v has exactly Y as neighborhood in X and the parent of v is adjacent to x (see Lines 4 to 18). To do this in $O(kn)$ time, the algorithm proceeds as follows: First, it computes in Line 5 the set P which contains all parents of interesting bottommost leaves. Clearly, this can be done in linear time. Next, the algorithm processes the vertices in P . Observe that further vertices might be added to P (see Line 18) during this processing. Let u be the currently processed vertex of P , let v be its child vertex, and let Y be the neighborhood of v in X . For each neighbor $x \in N(u) \cap X$, the algorithm checks whether there are already k augmenting paths between x and Y with a table lookup in Tab (see Line 10). If not, then the table entry is incremented by one (see Line 11) since u and v provide another augmenting path. If yes, then the edge $\{x, u\}$ is deleted in Line 13 (we show below that this does not change the maximum matching size). If u has degree two after processing all neighbors of u in X , then, by applying Reduction Rule 2, we can remove u and merge its two neighbors v and w . It follows from Lemma 7 that this application of Reduction Rule 2 can be done in $O(k)$ time. Hence, one iteration of the while loop requires $O(k)$ time and thus Algorithm 2 runs in $O(kn)$ time.

Recall that all vertices in $G - X$ that are free wrt. M_{G-X} are leaves. Thus, the changes to M_{G-X} by applying Reduction Rule 2 in Line 15 are as follows: First, the edge $\{u, v\}$ is removed and second the edge $\{w, q\}$ is replaced by $\{vw, q\}$ for some $q \in V$. Hence, the matching M_{G-X} after running Algorithm 2 has still at most k^2 free vertices and all of them are leaves.

It remains to prove that

- (a) the deletion of the edge $\{x, u\}$ in Line 13 results in an equivalent instance and
- (b) that the resulting instance has at most $2k^2(2^k + 1)$ bottommost leaves.

First, we show (a). To this end, assume towards a contradiction that the new graph $G' := G - \{x, u\}$ has a smaller maximum matching than G (clearly, G' cannot have a larger maximum matching). Thus, any maximum matching M_G for G has

to contain the edge $\{x, u\}$. This implies that the child v of u in $G - X$ is matched in M_G with one of its neighbors (except u): If v is free wrt. M_G , then deleting $\{x, u\}$ from M_G and adding $\{v, u\}$ yields another maximum matching not containing $\{x, u\}$, a contradiction. Recall that $N(v) = \{u\} \cup Y$ where $Y \subseteq X$ since v is a leaf in $G - X$. Thus, each maximum matching M_G for G contains for some $y \in Y$ the edge $\{v, y\}$. Observe that Algorithm 2 deletes $\{x, u\}$ only if there are at least k other interesting bottommost leaves v_1, \dots, v_k in $G - X$ such that their respective parent is adjacent to x and $N(v_i) \cap X = Y$ (see Lines 9 to 13). Since $|Y| \leq k$, it follows by the pigeonhole principle that at least one of these vertices, say v_i , is not matched to any vertex in Y . Thus, since v_i is an interesting bottommost leaf, it is matched to its only remaining neighbor: its parent u_i in $G - X$. This implies that there is another maximum matching

$$M'_G := (M_G \setminus \{\{v, y\}, \{x, u\}, \{u_i, v_i\}\}) \cup \{\{v_i, y\}, \{x, u_i\}, \{u, v\}\},$$

a contradiction to the assumption that all maximum matchings for G have to contain $\{x, u\}$.

We next show (b) that the resulting instance has at most $2k^2(2^k + 1)$ bottommost leaves. To this end, recall that there are at most $2k^2$ bottommost leaves that are not interesting (see discussion at the beginning of this subsection). Hence, it remains to upper-bound the number of interesting bottommost leaves. Observe that each parent u of an interesting bottommost leaf has to be adjacent to a vertex in X since otherwise u would have been deleted in Line 15. Furthermore, after running Algorithm 2, each vertex $x \in X$ is adjacent to at most $k2^k$ parents of interesting bottommost leaves (see Lines 10 to 13). Thus, the number of interesting bottommost leaves is at most k^22^k . Hence, the number of bottommost leaves is upper-bounded by $2k^2(2^k + 1)$. □

3.2.3 Step 6

In this subsection, we provide the final step of our kernelization algorithm. Recall that in the previous steps we have upper-bounded the number of bottommost leaves in $G - X$ by $O(k^22^k)$. We also computed a maximum matching M_{G-X} for $G - X$ such that at most k^2 vertices are free wrt. M_{G-X} and all free vertices are leaves in $G - X$. Using this, we next show how to reduce G to a graph of size $O(k^32^k)$. To this end we need some further notation. A leaf in $G - X$ that is not bottommost is called a *pendant*. We define T to be the *pendant-free tree (forest) of $G - X$* , that is, the tree (forest) obtained from $G - X$ by removing all pendants. The next observation shows that $G - X$ is not much larger than T . This allows us to restrict ourselves on giving an upper bound on the size of T instead of $G - X$.

Observation 5 Let $G - X$ be as described above with vertex set $V \setminus X$ and let T be the pendant-free tree (forest) of $G - X$ with vertex set V_T . Then, $|V \setminus X| \leq 2|V_T| + k^2$.

Proof Observe that $V \setminus X$ is the union of all pendants in $G - X$ and V_T . Thus, it suffices to show that $G - X$ contains at most $|V_T| + k^2$ pendants. To this end, recall that we have a maximum matching for $G - X$ with at most k^2 free leaves. Thus, there are at most k^2 leaves in $G - X$ that have a sibling which is also a leaf since from two leaves with the same parent at most one can be matched. Hence, all but at most k^2 pendants in $G - X$ have pairwise different parent vertices. Since all these parent vertices are in V_T , it follows that the number of pendants in $G - X$ is $|V_T| + k^2$. \square

We use the following observation to provide an upper bound on the number of leaves of T .

Observation 6 Let F be a forest, let F' be the pendant-free forest of F , and let B be the set of all bottommost leaves in F . Then, the set of leaves in F' is exactly B .

Proof First observe that each bottommost leaf of F is a leaf of F' since no bottommost leaf is removed and F' is a subgraph of F . Thus, it remains to show that each leaf v in F' is a bottommost leaf in F .

We distinguish two cases of whether or not v is a leaf in F : First, assume that v is not a leaf in F . Thus, all of its child vertices have been removed. Since we only remove pendants to obtain F' from F and since each pendant is a leaf, it follows that v is in F the parent of one or more leaves u_1, \dots, u_ℓ . Thus, by definition, all these leaves u_1, \dots, u_ℓ are bottommost leaves, a contradiction to the fact that they were deleted when creating F' .

Second, assume that v is a leaf in F . If v is a bottommost leaf, then we are done. Thus, assume that v is not a bottommost leaf and hence a pendant. However, since we remove all pendants to obtain F' from F , it follows that v is not contained in F' , a contradiction. \square

From Observation 6 it follows that the set B of bottommost leaves in $G - X$ is exactly the set of leaves in T . In the previous step we reduced the graph such that $|B| \leq 2k^2(2^k + 1)$ (see Proposition 2). Thus, T has at most $2k^2(2^k + 1)$ vertices of degree one and, since T is a tree (a forest), T also has at most $2k^2(2^k + 1)$ vertices of degree at least three. Let V_T^2 be the vertices of degree two in T and let $V_T^{\neq 2}$ be the remaining vertices in T . From the above it follows that $|V_T^{\neq 2}| \leq 4k^2(2^k + 1)$. Hence, it remains to upper-bound the size of V_T^2 . To this end, we will upper-bound the degree of each vertex in X by $O(k^2 2^k)$ and then use Reduction Rule 1 and 2. We will check for each edge $\{x, v\} \in E$ with $x \in X$ and $v \in V \setminus X$ whether we “need” it. This check will use the idea from the previous subsection where each vertex in X needs to reach each subset $Y \subseteq X$ at most k times via an augmenting path. Similarly as in the previous subsection, we want to keep “enough” of these augmenting paths. However, this time the augmenting paths might be long and different augmenting paths might overlap. To still use the basic approach, we use the following lemma stating that we can still somehow replace augmenting paths.

Lemma 8 *Let M_{G-X} be a maximum matching in the forest $G - X$. Let P_{uv} be an augmenting path for M_{G-X} in G from u to v . Let P_{wx} , P_{wy} , and P_{wz} be three internally vertex-disjoint augmenting paths from w to x , y , and z , respectively, such that P_{uv} intersects all of them. Then, there exist two vertex-disjoint augmenting paths with endpoints u , v , w , and one of the three vertices x , y , and z .*

Proof Label the vertices in P_{uv} alternating as *odd* or *even* with respect to P_{uv} so that no two consecutive vertices have the same label, u is *odd*, and v is *even*. Analogously, label the vertices in P_{wx} , P_{wy} , and P_{wz} as *odd* and *even* with respect to P_{wx} , P_{wy} , and P_{wz} , respectively, so that w is always *odd*. Since all these paths are augmenting, it follows that each edge from an even vertex to its succeeding odd vertex is in the matching M_{G-X} and each edge from an odd vertex to its succeeding even vertex is not in the matching. Observe that P_{uv} intersects each of the other paths at least at two consecutive vertices, since every second edge must be an edge in M_{G-X} . Since $G - X$ is a forest and all vertices in X are free with respect to M_{G-X} , it follows that the intersection of two augmenting paths is connected and thus a path. Since P_{uv} intersects the three augmenting paths from w , it follows that at least two of these paths, say P_{wx} and P_{wy} , have a “fitting parity”, that is, in the intersections of P_{uv} with P_{wx} and with P_{wy} the even vertices with respect to P_{uv} are either even or odd with respect to *both* P_{wx} and P_{wy} .

Assume without loss of generality that in the intersections of the paths the vertices have the same label with respect to the three paths (if the labels differ, then revert the ordering of the vertices in P_{uv} , that is, exchange the names of u and v and change all labels on P_{uv} to their opposite). Denote with v_s^1 and v_t^1 the first and the last vertex in the intersection of P_{uv} and P_{wx} . Analogously, denote with v_s^2 and v_t^2 the first and the last vertex in the intersection of P_{uv} and P_{wy} . Assume without loss of generality that P_{uv} intersects first with P_{wx} and then with P_{wy} . Observe that v_s^1 and v_s^2 are even vertices and v_t^1 and v_t^2 are odd vertices since the intersections have to start and end with edges in M_{G-X} (see Fig. 2 for an illustration). For an arbitrary path P and for two arbitrary vertices p_1, p_2 of P , denote by $p_1 - P - p_2$ the subpath of P from p_1 to p_2 . Observe that $u - P_{uv} - v_t^1 - P_{wx} - x$ and $w - P_{wy} - v_t^2 - P_{uv} - v$ are vertex-disjoint augmenting paths. \square

Algorithm description. We now provide the algorithm for Step 6 (see Algorithm 3 for pseudocode). The algorithm uses the same preprocessing (see Lines 1 to 3) as Algorithm 2. Thus, the algorithm can determine whether two vertices have the same neighborhood in X in constant time. As in Algorithm 2, Algorithm 3 uses a table Tab which has an entry for each vertex $x \in X$ and each set $Y \subseteq X$. The table is filled in such a way that the algorithm detected for each $y \in Y$ at least $\text{Tab}[x, Y]$ internally vertex-disjoint augmenting paths from x to y . The main part of the algorithm is the boolean function ‘Keep-Edge’ in Line 13 to 22 which makes the decision on whether to delete an edge $\{x, v\}$ for $v \in V \setminus X$ and $x \in X$. The function works as follows for edge $\{x, v\}$: Starting at v the graph will be explored along possible augmenting paths until a “reason” for keeping




Fig. 2 The situation in the proof of Lemma 8. The augmenting path from u to v intersects the two augmenting paths P_{wx} and P_{wy} from w to x and y , respectively. Bold edges indicate edges in the matching, dashed edges indicate odd-length alternating paths starting with the first and last edge not being in the matching. The gray paths in the background highlight the different augmenting paths: the initial paths from u to v , w to x , and x to y as well as the new paths from u to x and w to v as postulated by Lemma 8




Fig. 3 Illustration of the graph exploration of the function Keep-Edge in Algorithm 3: The vertices x and y are vertices in the feedback vertex set X . The vertices v_1, \dots, v_{10} are part of $G - X$ where v_{10} is a free leaf. The matching M_{G-X} is denoted by the thick edges. Three alternating paths are highlighted; each path represents an exploration of Keep-Edge from x that returns true: First, the path via v_3 ends in v_1 —a vertex with degree more than two in $G - X$ (see Line 14). The second path via v_4 ends in v_5 —a vertex connected to two vertices in X (here we assume that there are less than $6k^2$ paths from x to vertices adjacent to y and z ; see Lines 17–19). The third path via v_8 ends in the free leaf v_{10} (see Line 14)

the edge $\{x, v\}$ is found or no further exploration is possible (see Fig. 3 for an illustration).

Algorithm 3: An algorithm for computing Step 6 in $O(kn)$ time.

Input: A matching instance $(G = (V, E), s)$, a feedback vertex set $X \subseteq V$ of size k for G with $k < \log n$ and at most $k^2(2^k + 1)$ bottommost leaves in $G - X$, and a maximum matching M_{G-X} for $G - X$ with at most k^2 free vertices in $G - X$ that are all leaves.

Output: An equivalent matching instance (G', s') such that G' contains at most $O(k^3 2^k)$ vertices and edges.

```

1 Fix an arbitrary bijection  $f: 2^X \rightarrow \{1, \dots, 2^k\}$ 
2 foreach  $v \in V \setminus X$  do
3    $\left[ \begin{array}{l} \text{Set } f_X(v) \leftarrow f(N(v) \cap X) \quad // \text{ The number } f_X(v) < n \text{ can be read in constant} \\ \text{time.} \end{array} \right.$ 
4 Initialize a table Tab of size  $k \cdot 2^k$  with  $\text{Tab}[x, f(Y)] \leftarrow 0$  for  $x \in X, \emptyset \subsetneq Y \subseteq X$ 
5  $T \leftarrow$  pendant-free tree (forest) of  $G - X$ 
6  $V_T^{\geq 3} \leftarrow$  vertices in  $T$  with degree  $\geq 3$ 
7 foreach  $x \in X$  do
8   foreach  $v \in N(x) \setminus X$  do
9      $\left[ \begin{array}{l} \text{if } \text{Keep-Edge}(x, v) = \text{false} \text{ then} \quad // \text{ Is } \{x, v\} \text{ needed for an augmenting} \\ \text{path?} \\ \quad \left[ \begin{array}{l} \text{delete } \{x, v\} \end{array} \right. \end{array} \right.$ 
11 Exhaustively apply first Reduction Rule 1 and then Reduction Rule 2
12 return  $(G, s)$ .
13 Function  $\text{Keep-Edge}(x \in X, v \in V \setminus X)$ 
14   if  $v$  is free wrt.  $M_{G-X}$  or  $v \in V_T^{\geq 3}$  then return true
15    $w \leftarrow$  matched neighbor of  $v$  in  $M_{G-X}$ 
16   if  $w \in V_T^{\geq 3}$  or  $w$  is adjacent to a free leaf in  $G - X$  then return true
17   if  $w$  has at least one neighbor in  $X$  and  $\text{Tab}[x, f_X(w)] < 6k^2$  then
18      $\left[ \begin{array}{l} \text{Tab}[x, f_X(w)] \leftarrow \text{Tab}[x, f_X(w)] + 1 \\ \text{return true} \end{array} \right.$ 
20   foreach  $u \in N(w) \setminus \{v\}$  that is matched wrt.  $M_{G-X}$  and fulfills  $\{u, x\} \notin E$  do
21      $\left[ \begin{array}{l} \text{if } \text{Keep-Edge}(u, x) = \text{true} \text{ then return true} \end{array} \right.$ 
22   return false

```

If the vertex v is free wrt. M_{G-X} , then $\{x, v\}$ is an augmenting path and we keep $\{x, v\}$ (see Line 14). Observe that in Step 4 (see Proposition 1) we upper-bounded the number of free vertices by k^2 and all these vertices are leaves. Thus, we keep a bounded number of edges incident to x because the corresponding augmenting paths can end at a free leaf. We provide the exact bound below when discussing the size of the graph returned by Algorithm 3. In Line 14, the algorithm also stops exploring the graph and keeps the edge $\{x, v\}$ if v has degree at least three in T . The reason is to keep the graph exploration simple by following only degree-two vertices in T . This ensures that the running time for exploring the graph from x does not exceed $O(n)$. Since the number of vertices in T with degree at least three is bounded (see discussion after Observation 6), it follows that only a bounded number of such edges $\{x, v\}$ are kept.

If v is not free wrt. M_{G-X} , then it is matched with some vertex w . If w is adjacent to some leaf u in $G - X$ that is free wrt. M_{G-X} , then the path x, v, w, u is an augmenting path. Thus, the algorithm keeps in this case the edge $\{x, v\}$, see Line

16. Again, since the number of free leaves is bounded, only a bounded number of edges incident to x will be kept. If w has degree at least three in T , then the algorithm stops the graph exploration here and keeps the edge $\{x, v\}$, see Line 16. Again, this is to keep the running time at $O(kn)$ overall.

Let $Y \subseteq X$ denote the neighborhood of w in X . Thus the partial augmenting path x, v, w can be extended to each vertex in Y . Thus, if the algorithm did not yet find $6k^2$ paths from x to vertices whose neighborhood in X is also Y , then the table entry $\text{Tab}[x, f_X(w)]$ (where $f_X(w)$ encodes the set $Y = N(w) \cap X$) is increased by one and the edge $\{x, v\}$ will be kept (see Lines 18 and 19). (Here we need $6k^2$ paths since these paths might be long and intersect with many other augmenting paths, see proof of Proposition 3 for the details of why $6k^2$ is enough.) If the algorithm already found $6k^2$ “augmenting paths” from x to Y , then the neighborhood of w in X is irrelevant for x and the algorithm continues.

In Line 20, all above discussed cases to keep the edge $\{x, v\}$ do not apply and the algorithm extends the partial augmenting part x, v, w by considering the neighbors of w except v . Since the algorithm dealt with possible extensions to vertices in X in Lines 17 and 19 and with extensions to free vertices in $G - X$ in Line 14, it follows that the next vertex on this path has to be a vertex u that is matched wrt. M_{G-X} . Furthermore, since we want to extend a partial augmenting path from x , we require that u is not adjacent to x : otherwise the length-one path x, u would be another, shorter partial augmenting path from x to u and we do not need the currently stored partial augmenting path.

Statements on Algorithm 3. To show that Algorithm 3 indeed performs Step 6, we need further lemmas. For each edge $\{x, z\}$ with $x \in X$ and $z \in V \setminus X$ we denote by $P(x, z)$ the induced subgraph of $G - X$ on the vertices that are explored in the function `Keep-Edge` when called in Line 9 with x and z . More precisely, we initialize $P(x, z) := \emptyset$. Whenever the algorithm reaches Line 14, we add v to $P(x, z)$. Furthermore, whenever the algorithm reaches Line 17, we add w to $P(x, z)$. Similarly, when the recursive call in Line 21 returns true, then we add u to $P(x, z)$ in the recursive call (with u taking the role of v).

We next show that $P(x, z)$ is a path with at most one additional pendant.

Lemma 9 *Let $x \in X$ and $z \in V \setminus X$ be two vertices such that $\{x, z\} \in E$. Then, $P(x, z)$ is either a path or a tree with exactly one vertex z' having more than two neighbors in $P(x, z)$. Furthermore, z' has degree exactly three and z is a neighbor of z' .*

Proof We first show that all vertices in $P(x, z)$ except z and its neighbor z' have degree at most two in $P(x, z)$. Observe that having more vertices than z and z' in $P(x, z)$ requires Algorithm 3 to reach Line 20.

Let w be the currently last vertex when Algorithm 3 continues the graph exploration in Line 20. Observe that the algorithm therefore dealt with the case that w has degree at least three in the pendant-free tree T in Line 16. Thus, w is either a pendant leaf in $G - X$ or $w \notin V_T^{\geq 3}$ (that is, w has degree at most two in T). In the first case, there is no candidate to continue and the graph exploration stops. In the second case, w has degree at most two in T .

We next show that any candidate u for continuing the graph exploration in Line 21 is not a leaf in $G - X$. Assume toward a contradiction that u is a leaf in $G - X$. Since the parent w of u is matched with some vertex $v \neq u$ (this is how w is chosen, see Line 15), it follows that u is not matched. This implies that the function ‘Keep-Edge’ would have returned true in Line 16 and would not have reached Line 20, a contradiction. Thus, the graph exploration follows only vertices in T . Furthermore, the above argumentation implies that w is not adjacent to a leaf unless this leaf is its predecessor v in the graph exploration.

We now have two cases: Either w is not adjacent to a leaf in $G - X$ or $v = z$ is a leaf and $w = z'$ is its matched neighbor. In the first case, w has at most one neighbor $u \neq v$ since $w \notin V_T^{\geq 3}$. Hence, w has degree two in $P(x, z)$. In the second case, $w = z'$ has at most two neighbors $u \neq v$ and $u' \neq v$. Thus, z' has degree at most three. □

For $x \in X$ let

$$\mathcal{P}_x := \{P(x, v) \mid \{x, v\} \in E \wedge v \in V \setminus X\}$$

be the union of all induced subgraphs that Algorithm 3 explores from x .

Lemma 10 *There exists a partition of \mathcal{P}_x into $\mathcal{P}_x = \mathcal{P}_x^A \cup \mathcal{P}_x^B$ such that all graphs within \mathcal{P}_x^A and within \mathcal{P}_x^B are pairwise disjoint.*

Proof Since $G - X$ is a tree (or forest), $G - X$ is also bipartite. Let A and B be its two color classes (so $A \cup B = V \setminus X$). We define the two parts \mathcal{P}_x^A and \mathcal{P}_x^B as follows: A subgraph $P \in \mathcal{P}_x$ is in \mathcal{P}_x^A if the neighbor v of x in P is contained in A , otherwise P is in \mathcal{P}_x^B .

We show that all subgraphs in \mathcal{P}_x^A and \mathcal{P}_x^B are pairwise vertex-disjoint. To this end, assume toward a contradiction that two graphs $P, Q \in \mathcal{P}_x^A$ share some vertex. (The case $P, Q \in \mathcal{P}_x^B$ is completely analogous.) Let p_1 and q_1 be the first vertex in P and Q respectively, that is, p_1 and q_1 are adjacent to x in G . Observe that $p_1 \neq q_1$. Let $u \neq x$ be the first vertex that is in P and in Q . By Lemma 9, P and Q are paths or trees with at most one vertex of degree more than two and this vertex has degree three and is the neighbor of p_1 or q_1 , respectively. This implies together with $q_1, p_1 \in A$ that either $u = p_1$ or $u = q_1$. Assume without loss of generality that $u = p_1$. Since $p_1 \in A$ and $q_1 \in A$ and u is a vertex in Q , it follows that Algorithm 3 followed u in the graph exploration from q_1 in Line 21. However, this is a contradiction since the algorithm checks in Line 20 whether the new vertex u in the path is not adjacent to x . Thus, all subgraphs in \mathcal{P}_x^A and \mathcal{P}_x^B are pairwise vertex-disjoint. □

We next show that if $\text{Tab}[x, f(Y)] = 6k^2$ for some $x \in X$ and $Y \subseteq X$ (recall that f maps Y to a number, see Line 1), then there exist at least $3k^2$ internally vertex-disjoint augmenting paths from x to Y .

Lemma 11 *If in Line 17 of Algorithm 3 it holds for $x \in X$ and $Y \subseteq X$ that $\text{Tab}[x, f(Y)] = 6k^2$, then there exist in G wrt. M_{G-X} at least $3k^2$ alternating*

paths from x to vertices v_1, \dots, v_{3k^2} such that all these paths are pairwise vertex-disjoint (except x) and $N(v_i) \cap X = N(w) \cap X$ for all $i \in [3k^2]$.

Proof Note that each time $\text{Tab}[x, f(Y)]$ is increased by one (see Line 18), the algorithm found a vertex w such that there is an alternating path P from x to w and $N(w) \cap X = Y$. Furthermore, since the function Keep-Edge returns true in this case, the edge from x to its neighbor on P is not deleted in Line 10. Thus, there exist at least $6k^2$ alternating paths from x to vertices whose neighborhood in X is exactly Y . By Lemma 10, it follows that at least half of these $6k^2$ paths are vertex-disjoint. \square

The next lemma shows that Algorithm 3 is correct and runs in $O(kn)$ time.

Proposition 3 *Let $(G = (V, E), s)$ be a matching instance, let $X \subseteq V$ be a feedback vertex set of size k with $k < \log n$ and at most $2k^2(2^k + 1)$ bottommost leaves in $G - X$, and let M_{G-X} be a maximum matching for $G - X$ with at most k^2 free vertices in $G - X$ that are all leaves. Then, Algorithm 3 computes in $O(kn)$ time an equivalent instance (G', s') of size $O(k^3 2^k)$.*

Proof We split the proof into three claims, one for the correctness of the algorithm, one for the returned kernel size, and one for the running time. \square

Claim 1 *The input instance (G, s) is a yes-instance if and only if the instance (G', s') produced by Algorithm 3 is a yes-instance.*

Proof Observe that the algorithm changes the input graph only in two lines: Lines 10 and 11. By Lemma 1, applying Reduction Rules 1 and 2 yields an equivalent instance. Thus, it remains to show that deleting the edges in Line 10 is correct, that is, it does not change the size of a maximum matching. To this end, observe that deleting edges does not increase the size of a maximum matching. Thus, we need to show that the size of the maximum matching does not decrease. Assume toward a contradiction that it does.

Let $\{x, v\}$ be the edge whose deletion decreased the maximum matching size. Redefine G to be the graph before the deletion of $\{x, v\}$ and G' to be the graph after the deletion of $\{x, v\}$. Recall that Algorithm 3 gets as additional input a maximum matching M_{G-X} for $G - X$. Let $M_G := M_G^{\max}(M_{G-X})$ be a maximum matching for G with the largest possible overlap with M_{G-X} and let $G^M := G(M_{G-X}, M_G) = (V, M_{G-X} \triangle M_G)$ (see Sect. 2). Since $\{x, v\} \in M_G \setminus M_{G-X}$ and x is free wrt. M_{G-X} , it follows that there is a path P in G^M with one endpoint being x .

Recall (see Sect. 2) that since P is a path in G^M it follows that P is an augmenting path for M_{G-X} . Since all vertices in X are free wrt. M_{G-X} , it follows that all vertices in P except the endpoints are in $V \setminus X$. Let z be the second endpoint of this path P . We call a vertex on P an even (odd) vertex if it has an even (odd) distance to x on P . (So x is an even vertex and v and z are odd vertices). Observe that v is the only odd

vertex in P adjacent to x : Otherwise there would be another augmenting path from x to z which only uses vertices from P . This would imply the existence of another maximum matching that does not use $\{x, v\}$, a contradiction.

Let u be the neighbor of z in P . Since no odd vertex on P except v is adjacent to x , it follows that the graph exploration in the function `Keep-Edge` starting from x and v in Line 9 either reached u or returned true before. If $z \in V \setminus X$, then in both cases, the function `Keep-Edge` would have returned true in Line 9 and Algorithm 3 would not have deleted $\{x, v\}$, a contradiction. Thus, assume that $z \in X$. Therefore, the function `Keep-Edge` considered the vertex u in Line 17 but did not keep the edge $\{x, v\}$. Thus, when considering u , it holds that $\text{Tab}[x, f_X(u)] = 6k^2$, where $f_X(u)$ encodes $Y := N(u) \cap X$ and $z \in Y$.

By Lemma 11, it follows that there are $3k^2$ pairwise vertex-disjoint (except x) alternating paths from x to vertices u_1, \dots, u_{3k^2} with $N(u_i) \cap X = Y$. Thus, there is a set \mathcal{Q} of $3k^2$ internally vertex-disjoint paths from x to y in G . If one of the paths $Q \in \mathcal{Q}$ does not intersect any path in G^M , then reverting the augmentation along P and augmenting along Q would result in another maximum matching not containing $\{x, v\}$, a contradiction. Thus, assume that each path in \mathcal{Q} intersects at least one path in G^M .

For each two paths $Q_1, Q_2 \in \mathcal{Q}$ that intersect the same path P' in G^M it holds that each further path P'' in G^M can intersect at most one of Q_1 and Q_2 : Assume toward a contradiction that P'' does intersect both Q_1 and Q_2 . Since no path in G^M except P contains x and z it follows that all intersections between the paths are within $G - X$. Since P' and P'' are vertex-disjoint and Q_1 and Q_2 are internally vertex-disjoint, it follows that there is a cycle in $G - X$, a contradiction to the fact that X is a feedback vertex set.

Since $3k^2 > 3k + k^2$, it follows from the pigeon hole principle that there is a path $P' \in G^M$ that intersects at least three paths $Q_1, Q_2, Q_3 \in \mathcal{Q}$ such that no further path in G^M intersects them. We can now apply Lemma 8 and obtain two vertex-disjoint augmenting paths Q and Q' . Thus, reverting the augmentation along P and P' and augmenting along Q and Q'' yields another maximum matching for G which does not contain $\{x, v\}$, a contradiction. \square

Claim 2 *The graph G' returned by Algorithm 3 has $O(k^3 2^k)$ vertices and edges.*

Proof We first show that each vertex $x \in X$ has degree $O(k^2 2^k)$ in G' . To this end, we need to count the number of neighbors $v \in N(x) \setminus X$ where the function `Keep-Edge` returns true in Line 9. By Lemma 9, the function `Keep-Edge` explores the graph along one or two paths (essentially growing from one starting point into two directions). Recall that \mathcal{P}_x denotes the subgraphs induced by the graph exploration of `Keep-Edge` for the neighbors of x . By Lemma 10 there is a partition of \mathcal{P}_x into \mathcal{P}_x^A and \mathcal{P}_x^B such that within each part the subgraphs are pairwise vertex-disjoint. We consider the two parts independently. We start with bounding the number of graphs in \mathcal{P}_x^A where the function ‘`Keep-Edge`’ returned true (the analysis is completely analogous for \mathcal{P}_x^B).

Since all explored subgraphs are disjoint and all free vertices in $G - X$ wrt. M_{G-X} are leaves, it follows that Algorithm 3 returned at most k^2 times true in Line 16 due to w being adjacent to a free leaf in $G - X$. Also, the algorithm returns at most k^2 times true in Line 14 due to v being free. Furthermore, the algorithm returns at most $6k^2 \cdot 2^k$ times true in Line 19. Finally, we show that the algorithm returns at most $8k^2 \cdot (2^k - 1)$ times true in Line 14 and 16, respectively. It follows from the discussion below Observation 6 that T , the pendent-free tree of $G - X$, has at most $2k^2(2^k + 1)$ leaves (denoted by V_T^1) and $2k^2(2^k + 1)$ vertices of degree at least three (denoted by $V_T^{\geq 3}$). Let V_T be the vertices of T . Since T is a tree (or forest), it has more vertices than edges and hence

$$\sum_{v \in V_T} \text{deg}_T(v) < 2|V_T|$$

which implies

$$\sum_{v \in V_T^{\geq 3}} \text{deg}_T(v) < 2 \cdot |V_T^{\geq 3}| + |V_T^1|.$$

Thus, Algorithm 3 returns at most $2 \cdot |V_T^{\geq 3}| + |V_T^1| < 6k^2(2^k + 1)$ times true in Line 16 due to w being a vertex in $V_T^{\geq 3}$. Also, Algorithm 3 returns at most $|V_T^{\geq 3}| \leq 2k^2(2^k + 1)$ times true in Line 14 due to v being a vertex in $V_T^{\geq 3}$.

Summarizing, considering the graph explorations in \mathcal{P}_x^A , Algorithm 3 returns at most

$$k^2 + k^2 + 6k^2 \cdot 2^k + 8k^2(2^k + 1) \in O(k^2 2^k)$$

times true in the function Keep-Edge. Analogously, considering the graph explorations in \mathcal{P}_x^A , Algorithm 3 also returned at most $O(k^2 2^k)$ times true. Hence, each vertex $x \in X$ has degree at most $O(k^2 2^k)$ in G' .

We now show that the exhaustive application of first Reduction Rule 1 and then Reduction Rule 2 indeed results in a kernel of the claimed size. To this end, denote with $V_{G'-X}^1$, $V_{G'-X}^2$, and $V_{G'-X}^{\geq 3}$ the vertices that have degree one, two, and at least three in $G' - X$. We have $|V_{G'-X}^1| \in O(k^3 2^k)$ since each vertex in X has degree at most $O(k^2 2^k)$ and G' is reduced wrt. Reduction Rule 1. Next, since $G' - X$ is a forest (or tree), we have $|V_{G'-X}^{\geq 3}| < |V_{G'-X}^1|$ and thus $|V_{G'-X}^{\geq 3}| \in O(k^3 2^k)$. Finally, each degree-two vertex in G' needs at least one neighbor of degree at least three since G' is reduced with respect to Reduction Rule 2. Thus, each vertex in $V_{G'-X}^2$ is either incident to a vertex in X or adjacent to one of the at most $O(k^2 2^k)$ vertices in $G' - X$ that have degree at least three. Thus, $|V_{G'-X}^2| \in O(k^3 2^k)$. Summarizing, G' contains at most $O(k^3 2^k)$ vertices and edges. \square

Claim 3 Algorithm 3 runs in $O(kn)$ time.

Proof First, observe that Lines 1 to 6 can be done in $O(kn)$ time: The preprocessing and table initialization can be done in $O(kn)$ time as discussed in Sect. 3.2.2. Furthermore, T and $V_T^{\geq 3}$ can clearly be computed in $O(n + m) \leq O(kn)$ time. Second,

applying Reduction Rule 1 in $O(n + m)$ time is straightforward and Bartha and Kresz [2] showed how to apply Reduction Rule 2 in $O(n + m)$ time. Thus, it remains to show that each iteration of the foreach-loop in Line 7 can be done in $O(n)$ time.

By Lemma 10, the graphs \mathcal{P}_x explored from x can be partitioned into two parts such that within each part all subgraphs are vertex-disjoint. Thus, each vertex in $G - X$ is visited only twice during the execution of the function Keep-Edge. Furthermore, observe that in Line 17 and 18 the table can be accessed in constant time. Thus, the function Keep-Edge only checks once whether a vertex in $V \setminus X$ has a neighbor in X , namely in Line 20. This single check can be done in constant time. Since the rest of the computation is done on $G - X$ which has less than $|V \setminus X|$ edges, it follows that each iteration of the foreach-loop in Line 7 can indeed be done in $O(n)$ time.

This completes the proof of Proposition 3. \square

This completes the description of Step 6. Combining Steps 1 to 6 we obtain our kernelization algorithm for the parameter feedback vertex number.

Theorem 2 *MATCHING parameterized by the feedback vertex number k admits a kernel of size $2^{O(k)}$. It can be computed in $O(kn)$ time.*

Proof First, using the linear-time factor-four approximation of Bar-Yehuda et al. [1], we compute an approximate feedback vertex set X with $|X| \leq 4k$. Then, we apply Step 1 to 6 using Algorithms 1 to 3. By Propositions 1 to 3, this can be done in $O(kn)$ time and results in a kernel of size $O((4k)^3 2^{4k}) = 2^{O(k)}$. \square

Applying the $O(m\sqrt{n})$ -time algorithm for MATCHING [21] on the kernel yields the following.

Corollary 2 *MATCHING can be solved in $O(kn + 2^{O(k)})$ time, where k is the feedback vertex number.*

4 Kernelization for Matching on Bipartite Graphs

In this section, we investigate the possibility of efficient and effective preprocessing for BIPARTITE MATCHING. More specifically, we show a linear-time computable polynomial-size kernel with respect to the parameter distance to chain graphs. In the first part of this section, we provide the definition of chain graphs and describe how to compute the parameter. In the second part, we discuss the kernelization algorithm.

Definition and computation of the parameter. We first define chain graphs which are a subclass of bipartite graphs with special monotonicity properties.

Definition 1 [4] Let $G = (A, B, E)$ be a bipartite graph. Then G is a *chain graph* if each of its two color classes A, B admits a linear ordering wrt. neighborhood




Fig. 4 A chain graph. Note that the ordering $<_A$ of the vertices in A is going from left to right while the ordering $>_B$ of the vertices in B is going from right to left. The reason for these two orderings being drawn in different directions is that a maximum matching can be drawn as parallel edges, see e. g. the bold edges. In fact, Algorithm 4 computes such matchings with the matched edges being parallel to each other

inclusion, that is, $A = \{a_1, \dots, a_\alpha\}$ and $B = \{b_1, \dots, b_\beta\}$ where $N(a_i) \subseteq N(a_j)$ and $N(b_i) \subseteq N(b_j)$ whenever $i < j$.

Observe that if the graph G contains twins, then there is more than one linear ordering wrt. neighborhood inclusion. To avoid ambiguities, we fix for the vertices of the color class A (resp. B) in a chain graph $G = (A, B, E)$ one linear ordering $<_A$ (resp. $<_B$) such that, for two vertices $u, v \in A$ (resp. $u, v \in B$), if $u <_A v$ (resp. if $u <_B v$) then $N(u) \subseteq N(v)$. In the remainder of the section we consider a bipartite representation of a given chain graph $G = (A, B, E)$ where the vertices of A (resp. B) are ordered according to $<_A$ (resp. $<_B$) from left to right (resp. from right to left), as illustrated in Fig. 4.

For simplicity of notation we use in the following $<$ to denote the orderings $<_A$ and $<_B$ whenever the color class is clear from the context. Note that we use the direction left/right to indicate the ordering $<$. That is, for a vertex $a' \in A$ to the right (left) of $a \in A$ we have $a < a'$ ($a' > a$). In contrast, for a vertex $b' \in B$ to the right (left) of $b \in B$ we have $b > b'$ ($b < b'$).

We next show that we have a constant-factor approximation for the parameter and the corresponding vertex subset working in linear time. To this end, we use the following characterization of chain graphs. Here, $2K_2$ denote the one-regular graph on four vertices (with disjoint two edges).

Lemma 12 [4] *A bipartite graph is a chain graph if and only if it does not contain an induced $2K_2$.*

Lemma 13 *There is a linear-time factor-4 approximation for the problem of deleting a minimum number of vertices in a bipartite graph in order to obtain a chain graph.*

Proof Let $G = (A, B, E)$ be a bipartite graph. We compute a set $S \subseteq A \cup B$ such that $G - S$ is a chain graph and S is at most four times larger than a minimum size of

such a set. The algorithm iteratively tries to find a $2K_2$ and deletes the four corresponding vertices until no further $2K_2$ is found. Since in each $2K_2$, by Lemma 12, at least one vertex needs to be removed, the algorithm yields the claimed factor-4 approximation.

The details of the algorithm are as follows: First, it initializes $S = \emptyset$ and sorts the vertices in A and in B by their degree; the vertices in $A = \{a_1, \dots, a_\alpha\}$ in increasing order and the vertices in $B = \{b_1, \dots, b_\beta\}$ in decreasing order, that is, $\deg(a_1) \leq \dots \leq \deg(a_\alpha)$ and $\deg(b_1) \geq \dots \geq \deg(b_\beta)$. Since the degree of each vertex is at most $\max\{\alpha, \beta\}$, this can be done in linear time with e. g. Bucket Sort. At any stage the algorithm deletes all vertices of degree zero and all vertices which are adjacent to all vertices in the other partition. The deleted vertices are not added to S since these vertices cannot participate in a $2K_2$. Next, the algorithm recursively processes the vertices in A in a nondecreasing order of their degrees. Let $a \in A$ be a minimum-degree vertex and let $b \in B$ be a neighbor of a . Since b is not adjacent to all vertices in A (otherwise b would be deleted), there is a vertex $a' \in A$ that is not adjacent to b . Since $\deg(a) \leq \deg(a')$ it follows that a' has a neighbor b' that is not adjacent to a . Hence, the four vertices a, a', b, b' induce only two edges: $\{a, b\}$ and $\{a', b'\}$ and thus form a $2K_2$. Thus, the algorithm adds the four vertices to S , deletes them from the graph, and continues with a vertex in A that has minimum degree.

As to the running time, we now show that, after the initial sorting, the algorithm considers each edge only twice: Selecting a and b as described above can be done in $O(1)$ time. To select a' , the algorithm simply iterates over all vertices in A until it finds a vertex that is not adjacent to b . In this way at most $\deg(b) + 1$ vertices are considered. Similarly, by iterating over the neighbors of a' , one finds b' . Hence, the edges incident to a, a', b , and b' are used once to find the vertices and a second time when these vertices are deleted. Thus, using appropriate data structures, the algorithm runs in $O(n + m)$ time. \square

Kernelization overview. In the rest of this section, we provide a linear-time computable kernel for BIPARTITE MATCHING with respect to the parameter vertex deletion distance k to chain graphs. On a high level, our kernelization algorithm consists of two steps: First, we upper-bound by $O(k)$ the number of neighbors of each vertex in the deletion set. Second, we mark $O(k^2)$ special vertices and we use the monotonicity properties of chain graphs to upper-bound the number of vertices that lie between any two consecutive marked vertices, thus bounding the total size of the reduced graph to $O(k^3)$ vertices.

Step 1. Let $G = (A, B, E)$ be the bipartite input graph, where $V = A \cup B$, and let $X \subseteq V$ be a vertex subset such that $G - X$ is a chain graph. By Lemma 13, we can compute an approximate X in linear time. Our kernelization algorithm uses a specific maximum matching $M_{G-X} \subseteq E$ in $G - X$ with Algorithm 4 where all edges in M_{G-X} are “parallel” and all matched vertices are consecutive in the ordering \prec_A and \prec_B , see also Fig. 4. Since in convex graphs matching is linear-time solvable [25] and convex graphs are a superclass of chain graphs, this can be done in $O(n + m)$ time. We use M_{G-X} in our kernelization algorithm to obtain some local information about possible augmenting paths. For example, each augmenting path has at least one endpoint in X . Forming this into a data reduction rule, with s denoting the size of a maximum matching, yields the following.

Algorithm 4: An algorithm computing a maximum matching M in the chain graph G such that all edges in M are “parallel” (see Figure 4 for a visualization.)

Input: A chain graph $G = (V, E)$, $V = A \cup B$, $A = \{a_1, \dots, a_\alpha\}$ and $B = \{b_1, \dots, b_\beta\}$ with $N(a_i) \subseteq N(a_j)$ and $N(b_i) \subseteq N(b_j)$ for $i < j$.

Output: A maximum matching of G where all matched edges are parallel.

- 1 Compute the size s of a maximum matching in G using an algorithm of Steiner and Yeomans [25]
- 2 $M \leftarrow \{\{a_{\alpha-s+1}, b_\beta\}, \{a_{\alpha-s+2}, b_{\beta-1}\}, \dots, \{a_\alpha, b_{\beta-s+1}\}\}$
- 3 **return** M .

Reduction Rule 4 If $|M_{G-X}| \geq s$, then return a trivial yes-instance; if $s > |M_{G-X}| + k$, then return a trivial no-instance.

The correctness of Reduction Rule 4 follows from Observation 2.

We will show next that there is a maximum matching M_G for G in which each vertex in X is either matched with another vertex in X or with a “small-degree vertex” in $G - X$. This means that an augmenting path starting at some vertex in X will “enter” the chain graph $G - X$ in a small-degree vertex. We now formalize this concept. Recall that $u < v$ implies $N(u) \subseteq N(v)$. For a vertex $x \in X$ we define $N_{\text{small}}^{V \setminus X}(x)$ to be the set of the k neighbors of x in $V \setminus X$ with the smallest degree, formally,

$$N_{\text{small}}^{V \setminus X}(x) := \{w \in N(x) \setminus X \mid k > |\{u \in N(x) \setminus X \mid u < w\}|\}.$$

Lemma 14 Let $G = (V, E)$ be a bipartite graph and let $X \subseteq V$ be a vertex set such that $G - X$ is a chain graph. Then, there exists a maximum matching M_G for G such that every matched vertex $x \in X$ is matched to a vertex in $N_{\text{small}}^{V \setminus X}(x) \cup X$.

Proof Assume, towards a contradiction, that there is no such matching M_G . Let M'_G be a maximum matching for G that maximizes the number of vertices $x \in X$ that are matched to a vertex in $N_{\text{small}}^{V \setminus X}(x) \cup X$, that is, let M'_G maximize $|\{x \in X \mid \{u, x\} \in M'_G \wedge u \in N_{\text{small}}^{V \setminus X}(x) \cup X\}|$. Let $x \in X$ be a vertex that is not matched with any vertex in $N_{\text{small}}^{V \setminus X}(x) \cup X$, that is, x is matched to a vertex $u \in V \setminus (N_{\text{small}}^{V \setminus X}(x) \cup X)$. If there is an unmatched vertex $w \in N_{\text{small}}^{V \setminus X}(x)$ in M'_G , then the matching $M''_G := M'_G \cup \{\{x, w\}\} \setminus \{\{u, x\}\}$ is a maximum matching with more vertices $x \in X$ (compared to M'_G) that are matched to a vertex in $N_{\text{small}}^{V \setminus X}(x) \cup X$, a contradiction. Hence, assume that there is no free vertex in $N_{\text{small}}^{V \setminus X}(x)$. Since $|N_{\text{small}}^{V \setminus X}(x)| = |X| = k$, it follows that at least one vertex $w \in N_{\text{small}}^{V \setminus X}(x)$ is matched to a vertex $v \in V \setminus X$. Observe that, by definition of $N_{\text{small}}^{V \setminus X}(x)$, we have $N_{G-X}(w) \subseteq N_{G-X}(u)$. Thus, we have $\{u, v\} \in E$ and thus, $M''_G := M'_G \cup \{\{x, w\}, \{u, v\}\} \setminus \{\{u, x\}, \{w, v\}\}$ is a maximum matching with more vertices in X (compared to M'_G) fulfilling the condition of the lemma, a contradiction. □




Fig. 5 A chain graph with a maximum matching (thick edges) and two additional vertices x and y . An augmenting path $P = x, a_1, b_1, a_2, b_2, \dots, a_7, b_7, y$ with $a_1 < a_\ell$ implies that there is a shorter augmenting path $P' = x, a_1, b_1, a_7, b_7, y$ of length five (indicated by the gray background) in the input graph since $N(a_1) \subseteq N(a_7)$

Based on Lemma 14, we can provide our next data reduction rule.

Reduction Rule 5 Let (G, s) be an instance reduced with respect to Reduction Rule 4 and let $x \in X$. Then delete all edges between x and $V \setminus N_{\text{small}}^{V \setminus X}(x)$.

Clearly, Reduction Rule 5 can be exhaustively applied in $O(n + m)$ time by one iteration over A and B in the ordering \prec .

Step 2. For the second step of our kernelization algorithm, we first mark a set K of $O(k^2)$ vertices that are kept in the graph (and thus will end up in the kernel): Keep all vertices of X . For each vertex $x \in X$ keep all vertices in $N_{\text{small}}^{V \setminus X}(x)$ and if a kept vertex is matched wrt. M_{G-X} , then keep also the vertex with which it is matched. Formally, we have:

$$K := X \cup \{v \mid \exists x \in X : v \in N_{\text{small}}^{V \setminus X}(x) \vee (\{u, v\} \in M_{G-X} \wedge u \in N_{\text{small}}^{V \setminus X}(x))\}.$$

Observe that exhaustively applying Reduction Rule 5 ensures that K is of size at most $2k^2$.

Next, we use the monotonicity properties of the chain graph to show that it suffices to keep for each vertex $v \in K$ at most k vertices to the right and to the left of v . Consider an augmenting path $P = x, a_1, b_1, \dots, a_\ell, b_\ell, y$ from a vertex $x \in B \cap X$ to a vertex $y \in A \cap X$. Observe that if $a_1 < a_\ell$, then also $\{b_1, a_\ell\} \in E$ and thus $P' = x, a_1, b_1, a_\ell, b_\ell, y$ is an augmenting path (see Fig. 5 for a visualization). Furthermore, the vertices in the augmenting path P' are a subset of $K \cup X$ and, thus, by keeping these vertices (and the edges between them), we also keep the augmenting path P' in our kernel. Hence, it remains to consider the more complicated case that $a_\ell < a_1$ (see Fig. 6). To this end, we next show that in certain “areas” of the chain graph $G - X$ the number of augmenting paths “passing through” such an area is upper-bounded. To specify an “area”, we need the following definition.

Definition 2 Let $G = (A, B, E)$ be a chain graph and let M be a matching in G . Furthermore let $a \in A, b \in B$ with $\{a, b\} \in M$. Then $\#lmv(b, M)$ (resp. $\#rmv(a, M)$) is

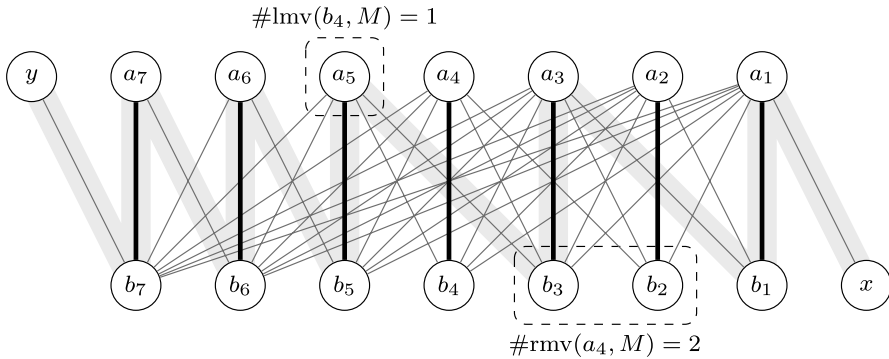


Fig. 6 The graph from Fig. 5 with the only difference being that the positions of x and y are exchanged (and the vertex names are adjusted accordingly, so x is still adjacent to a_1). Again, the thick, black edges denote a maximum matching M for the chain graph (containing all vertices except x and y). The graph contains an augmenting path $P = x, a_1, b_1, a_2, b_2, \dots, a_7, b_7, y$ with $a_1 > a_7$. In contrast to the example displayed in Fig. 5, there is no augmenting x - y -path of length five. The shortest augmenting x - y -path is displayed. For the edge $\{a_4, b_4\} \in M$ the vertices certifying that $\#lrmv(b_4, M) = 1$ and $\#rmv(a_4, M) = 2$ are highlighted by dashed boxes (see Definition 2)

the number of neighbors of b (resp. of a) that are to the left of a (resp. to the right of b); formally:

$$\#lrmv(b, M) := |\{a' \in N(b) \mid a' < a\}|, \#rmv(a, M) := |\{b' \in N(a) \mid b' < b\}|.$$

In Definition 2 the terms “left” and “right” refer to the ordering of the vertices of A and B in the bipartite representation of G , as illustrated in Fig. 6. The abbreviation $\#rmv$ ($\#lrmv$) stands for “number of vertices right (left) of the matched vertex”. We set

$$\#lrmv(a, M) := \#lrmv(b, M)$$

and

$$\#rmv(b, M) := \#rmv(a, M).$$

Finally, we define

$$\#rmv(a_1, a_2, M) := \min_{a_1 < a' < a_2} \{\#rmv(a', M)\}$$

for $a_1, a_2 \in A$ and

$$\#lrmv(b_1, b_2, M) := \min_{b_2 < b' < b_1} \{\#lrmv(b', M)\}$$

for $b_1, b_2 \in B$. For example, in the graph displayed in Fig. 6, we have $\#lrmv(b_2, b_3, M) = 2$ and $\#lrmv(b_3, b_6, M) = 1$.

With these definitions, we can show a limit on the number of augmenting paths that can “cross” an edge in M_{G-X} .

Lemma 15 *Let $G = (A, B, E)$ be a chain graph and M be a maximum matching for G computed by Algorithm 4. Let $a, b \in V$ with $\{a, b\} \in M$. Then the number of vertex-disjoint alternating paths that (1) start and end with edges not in M and that (2) have endpoints left of a and right of b is at most $\min\{\#l_{mv}(b, M), \#r_{mv}(a, M)\}$.*

Proof We prove the case $\#l_{mv}(b, M) \leq \#r_{mv}(a, M)$, that is, $\min\{\#l_{mv}(b, M), \#r_{mv}(a, M)\} = \#l_{mv}(b, M)$. The case $\#l_{mv}(b, M) > \#r_{mv}(a, M)$ follows by symmetry (with switched roles of a and b). Let $\#aug$ denote the number of vertex-disjoint alternating paths from $\{a' \in A \mid a' < a\}$ to $\{b' \in B \mid b' < b\}$ such that the first and last edge are not in M (see Fig. 6 for an example with $\#aug = 1$ for $a = a_4$ and $b = b_4$). Furthermore, let $a_1^b, \dots, a_{\#l_{mv}(b, M)}^b$ be the neighbors of b that are to the left of a , that is, $a_1^b < a_2^b < \dots < a_{\#l_{mv}(b, M)}^b < a$. Since G is a chain graph it follows that no vertex $a' \in A$ with $a' < a_1^b$ is adjacent to any vertex $b' \in B$ with $b' \leq b$. Furthermore, for any edge $\{a', b'\} \in E$ with $a < a'$ and $b < b'$ it follows from the construction of M (see Algorithm 4) that $\{a', b'\} \notin M$. Hence, any of these alternating paths has to contain at least one vertex from $a_1^b, \dots, a_{\#l_{mv}(b, M)}^b$. Since the alternating paths are vertex-disjoint it follows that $\#aug \leq \#l_{mv}(b, M)$. \square

From the previous lemma, we directly obtain the following.

Lemma 16 *Let $G = (A, B, E)$ be a chain graph and let M be the maximum matching for G computed by Algorithm 4. Let $a_1, a_2 \in A$ and $b_1, b_2 \in B$ with $\{a_1, b_1\}, \{a_2, b_2\} \in M$ with $a_1 < a_2$. Then there are at most $\#l_{mv}(b_1, b_2, M)$ vertex-disjoint alternating paths that (1) start and end with edges not in M and that (2) have endpoints left of a_1 and right of b_2 .*

Lemma 16 states that the number of augmenting paths passing through the “area” between a_1 and a_2 is bounded. Using this, we want to replace this area by a gadget with $O(k)$ vertices. To this end, we need further notation. For each kept vertex $v \in K$, we may also keep some bounded number of vertices to the right and to the left of v . We call these vertices the *left buffer (right buffer)* of v .

Definition 3 Let $G = (A, B, E)$ be a chain graph and let M be the maximum matching for G computed by Algorithm 4. Let $a_1, a_2 \in A$ and $b_1, b_2 \in B$ with $\{a_1, b_1\}, \{a_2, b_2\} \in M$ and $a_1 < a_2$. Then the (at most) $\#l_{mv}(b_1, b_2, M)$ vertices to the right of a_1 form the *right buffer* $B^r(a_1, M)$ of a_1 ; formally,

$$B^r(a_1, M) := \{a \in A \mid a_1 < a \wedge |\{a' \in A \mid a_1 < a' < a\}| \leq \min\{\#l_{mv}(b_1, b_2, M), k\}\}.$$

Analogously,

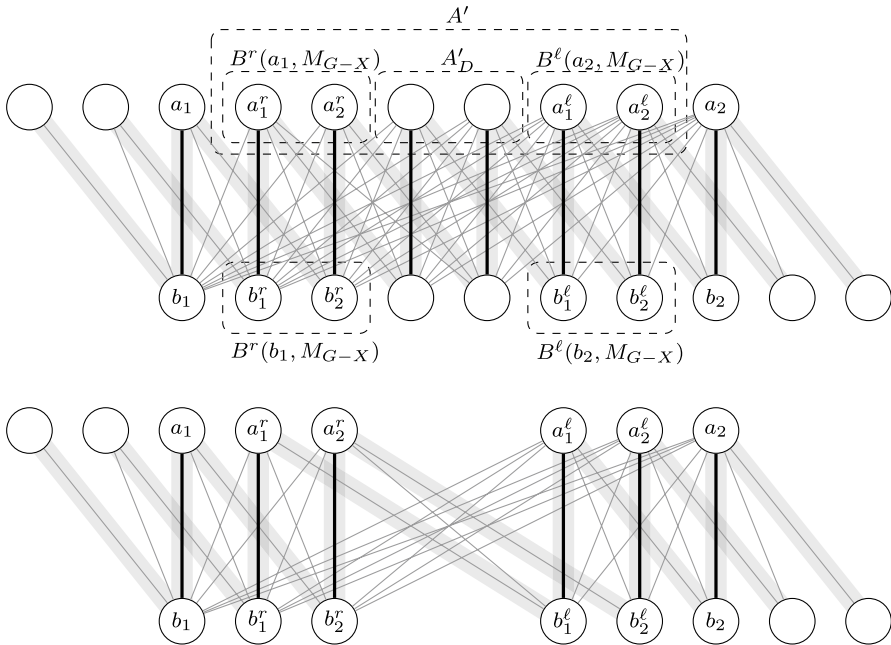


Fig. 7 An example for the application of Reduction Rule 6. Top: A part of a chain graph $G - X$ is displayed and the thick edges indicate parts of the maximum matching M_{G-X} . There are two vertex-disjoint augmenting paths from the two vertices left of a_1 to the two vertices right of b_2 highlighted by gray background. Moreover, $\#lmv(b_1, b_2, M_{G-X}) = 2$. Bottom: The part of the graph after applying Reduction Rule 6 and the corresponding two vertex-disjoint augmenting paths. Note that the only edges that are in the graph below but not above are the edges from vertices in $B^r(a_1, M_{G-X})$ to $B^l(b_2, M_{G-X})$

$$\begin{aligned}
 B^l(a_2, M) &:= \{a \in A \mid a < a_2 \wedge \\
 &\quad |\{a' \in A \mid a < a' < a_2\}| \leq \min\{\#lmv(b_1, b_2, M), k\}\}, \\
 B^r(b_1, M) &:= \{b \in B \mid b < b_1 \wedge \\
 &\quad |\{b' \in B \mid b < b' < b_1\}| \leq \min\{\#lmv(b_1, b_2, M), k\}\}, \\
 B^l(b_2, M) &:= \{b \in B \mid b_2 < b \wedge \\
 &\quad |\{b' \in B \mid b_2 < b' < b\}| \leq \min\{\#lmv(b_1, b_2, M), k\}\}.
 \end{aligned}$$

Note that in Definition 3 each of the sets $B^r(a_1, M)$, $B^l(a_2, M)$, $B^r(b_1, M)$, and $B^l(b_2, M)$ depends on all four vertices a_1, a_2, b_1, b_2 ; we omit these dependencies from the names for the sake of brevity.

The basic idea is now to delete vertices “outside” these buffers. See Fig. 7 for an illustrating example of the following data reduction rule formalizing this idea.

Reduction Rule 6 Let (G, s) be an instance reduced with respect to Reduction Rule 4. Let $a_1, a_2 \in K \cap A$ with $a_1 < a_2$ and $\{a_1, b_1\}, \{a_2, b_2\} \in M_{G-X}$ such that $A' := \{a \in A \mid a_1 < a < a_2\}$ is of size at least $2 \cdot \min\{\#lmv(b_1, b_2, M), k\} + 1$ and $A' \cap K = \emptyset$. Then delete all vertices in $A'_D := A' \setminus (B^r(a_1, M_{G-X}) \cup B^l(a_2, M_{G-X}))$

and their matched neighbors in B , add all edges between the vertices in the right buffer of a_1 and the vertices in the left buffer of b_2 , and decrease s by $|A'_D|$.

Lemma 17 *Reduction Rule 6 is correct and can be exhaustively applied in $O(n + m)$ time.*

Proof We first introduce some notation and provide some general observations. Then we show the correctness in two separate claims. Finally, we discuss the running time.

Let $a_1, a_2, b_1,$ and b_2 be as stated in Reduction Rule 6. Denote by A' (resp. B') the set of vertices between a_1 and a_2 (resp. between b_1 and b_2). Further denote by $A'_D \subseteq A'$ and $B'_D \subseteq B'$ the sets of deleted vertices. Note that $|A'| = |B'|$ and $|A'_D| = |B'_D|$ since M_{G-X} was produced by Algorithm 4. Denote the vertices in the buffers of $a_1, a_2, b_1,$ and b_2 by

$$B^x(y_z, M_{G-X}) := \{y_1^x, \dots, y_{\min\{\#l_{mv}(b_1, b_2, M_{G-X}), k\}}^x\}$$

for $x \in \{r, \ell\}, y \in \{a, b\}, z \in [2]$, and $x = r \iff z = 1$ (see Fig. 7 for examples of the concrete variable identifier).

Since the input instance is reduced with respect to Reduction Rule 4, it follows that $s - k \leq |M_{G-X}| < s$. Denote by $M_{G'-X} := M_{G-X} \cap E'$ the matching obtained from M_{G-X} by deleting all edges not in the reduced graph G' . Recall that $s' = s - |A'_D|$. We next show in Claims 4 and 5 that the input instance $(G = (V, E), s)$ is a yes-instance if and only if the produced instance $(G' = (V', E'), s')$ is a yes-instance. Before we present these two claims, observe that there is a perfect matching between the vertices in A'_D and B'_D , and thus

$$s - |M_{G-X}| = s' - |M_{G'-X}|. \tag{1}$$

Claim 4 *If (G, s) is a yes-instance, then (G', s') is a yes-instance.*

Proof Recall that $M_G^{\max}(M_{G-X})$ is a maximum matching for G minimizing the size of $M_{G-X} \triangle M_G^{\max}(M_{G-X})$. Since (G, s) is a yes-instance it holds that $|M_G^{\max}(M_{G-X})| \geq s$. For brevity we set $G^M := G(M_{G-X}, M_G^{\max}(M_{G-X})) = (V, M_{G-X} \triangle M_G^{\max}(M_{G-X}))$. Note that G^M is a graph that only contains odd-length paths (see Sect. 2). We will show that there are as many vertex-disjoint augmenting paths for M_{G-X} in G' as there are paths in G^M . This will show that G' contains a matching of size

$$\begin{aligned} |M_{G'-X}| + |M_G^{\max}(M_{G-X})| - |M_{G-X}| &\geq |M_{G'-X}| + s - |M_{G-X}| \\ &\stackrel{(1)}{=} |M_{G'-X}| + s' - |M_{G'-X}| = s'. \end{aligned}$$

To this end, observe that all paths that do not use vertices in $V \setminus V' = A'_D \cup B'_D$ are also contained in G' . Thus, consider the paths in G^M that use vertices in $V \setminus V'$. Denote by \mathcal{P}^M the set of all paths in G^M using vertices in $V \setminus V'$ and set $t := |\mathcal{P}^M|$. Consider now an arbitrary $i \in [t]$, and let $P_i^M \in \mathcal{P}^M$.

Denote by $v_1^i, v_2^i, \dots, v_{p_i}^i$ the vertices in P_i^M in the corresponding order, that is, v_1^i and $v_{p_i}^i$ are the endpoints of P_i^M and we have $\{v_{2j}^i, v_{2j+1}^i\} \in M_{G-X}$ for all $j \in [p_i/2]$. Observe that exactly one endpoint of P_i^M is in A and the other endpoint is in B , since P_i^M is an odd-length path. Assume without loss of generality that $v_1^i \in A$ and $v_{p_i}^i \in B$. Thus, the vertices in P_i^M with odd (even) index are in A (B).

We next show that for any two vertices v_j^i, v_ℓ^i of P_i^M with $j < \ell < p_i - 1$ and both being in $A \setminus X$, it follows that $v_j^i < v_\ell^i$. First, observe that if $j = 1$, then $v_j^i \in A \setminus X$ is a free vertex wrt. M_{G-X} . Since v_ℓ^i is matched wrt. M_{G-X} and since M_{G-X} is computed by Algorithm 4, it follows that $v_1^i < v_\ell^i$. Thus, assume that $j > 1$ and $\ell > 1$ (thus $j \geq 3$ and $j > 1$). Assume toward a contradiction that $v_\ell^i < v_j^i$. Since $\ell < p_i - 1$, we have $v_{\ell+1}^i \in B \setminus X$ and since $G - X$ is a chain graph, it follows that $\{v_j^i, v_{\ell+1}^i\} \in E$, a contradiction to Observation 1. Thus, $v_j^i < v_\ell^i$.

We next show that the path P_i^M contains at least one vertex v_j^i left of a_1 and at least one vertex v_ℓ^i right of b_2 . Recall that M_{G-X} was computed by Algorithm 4 and, thus, the free vertices are the smallest wrt. the ordering $<$ (see also Fig. 4). Thus, if one endpoint of P_i^M is in $(A \cup B) \setminus X$, then this vertex is either v_1^i and left of a_1 or it is $v_{p_i}^i$ and right of b_2 . Thus, assume that the endpoints of P_i^M are in X . We showed in the previous paragraph that $v_3^i < v_5^i < \dots < v_{p_i-3}^i$. Thus, we also have $v_{p_i-2}^i < v_{p_i-4}^i < \dots < v_2^i$ since M_{G-X} is computed by Algorithm 4. Since we assumed some vertices of P_i^M to be in $V \setminus V'$, it follows that for at least one vertex v_{2j+1}^i it holds that $a_1 < v_{2j+1}^i < a_2$. Furthermore since by assumption no vertex between a_1 and a_2 or between b_1 and b_2 is in K , it follows that $v_3^i < a_1$ (since $v_3^i \in K$) and $v_{p_i-2}^i < b_2$ (since $v_{p_i-2}^i \in K$).

For each $i \in [t]$ denote by $a_i^{P_i^M}$ the last vertex on the path P_i^M that is not right of a_1 , that is, $a_i^{P_i^M}$ is the vertex on P_i^M such that for each vertex $a' \in A \setminus X$ that is in P_i^M it holds that $a_1 < a'$ or $a' \leq a_i^{P_i^M}$. It follows from the previous paragraph that $a_i^{P_i^M}$ exists. Analogously to $a_i^{P_i^M}$, for each $i \in [t]$ denote by $b_i^{P_i^M}$ the first vertex on the path P_i^M that is not left of b_2 , that is, $b_i^{P_i^M}$ is the vertex on P_i^M such that for each vertex $b' \in B$ that is in P_i^M it holds that $b_2 < b'$ or $b' \leq b_i^{P_i^M}$. This means that in G^M there is for each $i \in [t]$ an alternating path from $a_i^{P_i^M}$ to $b_i^{P_i^M}$ starting and ending with non-matched edges and all these paths are pairwise vertex-disjoint. We show that also in G' there are pairwise vertex-disjoint alternating paths from $a_i^{P_i^M}$ to $b_i^{P_i^M}$. Assume without loss of generality that $a_1^{P_1^M} < a_2^{P_2^M} < \dots < a_t^{P_t^M}$. Since in each path $P_i^M, i \in [t]$, the successor of $a_i^{P_i^M}$ is to the right of b_1 , it follows that $a_i^{P_i^M}$ has at least i neighbors right of b_1 . Since the right buffer of b_1 contains the $\#lmv(b_1, b_2, M_{G-X}) \geq t$ (see Lemma 16) vertices to the right of b_1 , we have $\{a_i^{P_i^M}, b_i^r\} \in E$. By symmetry, we have $\{b_i^{P_i^M}, a_i^r\} \in E$. Recall that M_{G-X} forms a perfect matching between $B^r(b_1, M_{G-X})$ and $B^r(a_1, M_{G-X})$ as well as between $B^\ell(a_2, M_{G-X})$ and $B^\ell(b_2, M_{G-X})$. Since Reduction Rule 6 added all edges between $B^r(a_1, M_{G-X})$ and $B^\ell(b_2, M_{G-X})$ to E' , it follows that each path P_i^M can be completed as follows: $a_i^{P_i^M}, b_i^r, a_i^r, b_i^\ell, a_i^\ell, b_i^{P_i^M}$; note that exactly the edges $\{b_i^r, a_i^r\}$ and $\{b_i^\ell, a_i^\ell\}$ are in M_{G-X} . Thus, each path in \mathcal{P}^M can be replaced by an augmenting path for $M_{G'-X}$ in G' and all these augmenting paths are vertex-disjoint. Thus, there

are as many augmenting paths for $M_{G'-x}$ in G' as there are paths in G^M and therefore (G', s') is a yes-instance. □

Claim 5 *If (G', s') is a yes-instance, then (G, s) is a yes-instance.*

Proof Let $M_{G'}$ a maximum matching for G' . Observe that $|M_{G'}| \geq s'$. We construct a matching M_G for G as follows. First, copy all edges from $M_{G'} \cap E$ into M_G . Second, add all edges from $M_{G-x} \cap \binom{V \setminus V'}{2}$, that is, a perfect matching between A'_D and B'_D is added to M_G . Observe that if all edges in $M_{G'}$ are also in E , then M_G is a matching of size s in G . Thus, assume that some edges in $M_{G'}$ are not in E , that is, $\{a'_i, b'_i\}, \dots, \{a'_t, b'_t\} \in M_{G'} \setminus E$ for some $t \in [\#lmv(b_1, b_2, M_{G-x})]$. Observe that $s - |M_G| \leq t$. Clearly, the vertices $a'_1, \dots, a'_t, b'_1, \dots, b'_t$ are free with respect to M_G .

We show that there are t pairwise vertex-disjoint augmenting paths from $\{a'_1, \dots, a'_t\}$ to $\{b'_1, \dots, b'_t\}$; note, however, that these paths are not necessarily from a'_r to b'_r , where $r \in [t]$. To this end, recall that, by definition of $\#lmv(b_1, b_2, M_{G-x})$, each vertex $b \in B$ with $b_2 < b < b_1$ has at least $\#lmv(b_1, b_2, M_{G-x})$ neighbors to the left of its matched neighbor. This allows us to iteratively find augmenting paths as follows: To create the q^{th} augmenting path P_q start with some vertex b'_{j_q} . Denote by v the last vertex added to P_q (in the beginning we have $v = b'_{j_q}$). If $v \in A$, then add to P_q the neighbor matched to v . If $v \in B$, then do the following: if v is adjacent to a vertex $a \in \{a'_1, \dots, a'_t\}$, then add a to P_q , otherwise add the leftmost neighbor of v to P_q . Repeat this process until P_q contains a vertex from $\{a'_1, \dots, a'_t\}$. After we found P_q , remove all vertices of P_q from G . If $q < t$, then continue with P_{q+1} . Observe that any two vertices of P_q that are in A have at least $\#lmv(b_1, b_2, M_{G-x}) - 1$ other vertices of A between them (in the ordering of the vertices of A , see Fig. 4). Thus, after a finite number of steps, P_q will reach a vertex in $\{a'_1, \dots, a'_t\}$. Furthermore, it follows that after removing the vertices of P_q it holds that $\#lmv(b_1, b_2, M_{G-x})$ is decreased by exactly one: P_q contains for each vertex $b \in B$ at most one vertex among the $\#lmv(b_1, b_2, M_{G-x})$ neighbors of b that are directly to the left of its matched neighbor in M_{G-x} . Thus, in each iteration we have $\#lmv(b_1, b_2, M_{G-x}) > 0$. It follows that the above procedure constructs t vertex-disjoint augmenting paths from $\{a'_1, \dots, a'_t\}$ to $\{b'_1, \dots, b'_t\}$. Hence, G contains a matching of size s and thus (G, s) is a yes-instance. □

The correctness of the data reduction rule follows from the previous two claims. It remains to prove the running time. To this end, observe that the matching M_{G-x} is given. Computing all degrees of G can be done in $O(n + m)$ time. Also $\#lmv(v, M_{G-x})$ can be computed in linear time: For each vertex $b \in B$ one has to check for each neighbor of b whether it is to the left of b 's matched neighbor and to adjust $\#lmv(b, M_{G-x})$ accordingly. Furthermore, computing $\#lmv(b_1, b_2, M_{G-x})$ and removing the vertices in A'_D and B'_D can be done

in $O(\sum_{b \in B'} \deg(b) + \sum_{a \in A'} \deg(a))$ time. Thus, Reduction Rule 6 can be exhaustively applied in $O(n + m)$ time. \square

We next upper-bound the number of free vertices with respect to M_{G-X} . Let

$$A_{\text{free}} := \{a \in A \mid a \text{ is free with respect to } M_{G-X}\}$$

and

$$A_{\text{free}}^k := \{a \in A_{\text{free}} \mid |\{a' \in A_{\text{free}} \mid a \leq a'\}| \leq k\},$$

that is, A_{free}^k contains the k rightmost free vertices in $A \setminus X$. Observe that all vertices in A_{free}^k are left of M_{G-X} . Analogously, denote by B_{free}^k the set containing the k leftmost free vertices in $B \setminus X$.

Reduction Rule 7 Let (G, s) be an instance reduced with respect to Reduction Rule 4. Then delete all vertices in $(A_{\text{free}} \setminus (K \cup A_{\text{free}}^k)) \cup (B_{\text{free}} \setminus (K \cup B_{\text{free}}^k))$.

Lemma 18 *Reduction Rule 7 is correct and can be applied in $O(n + m)$ time.*

Proof The running time is clear. It remains to show the correctness. Let (G, s) be the input instance reduced with respect to Reduction Rule 4 and let (G', s) be the instance produced by Reduction Rule 7. We show that deleting the vertices in $A_{\text{free}} \setminus (K \cup A_{\text{free}}^k)$ yields an equivalent instance. It then follows from symmetry that deleting the vertices in $B_{\text{free}} \setminus (K \cup B_{\text{free}}^k)$ yields also an equivalent instance.

We first show that if (G, s) is a yes-instance, then also the produced instance (G', s) is a yes-instance. Let (G, s) be a yes-instance and M_G be a maximum matching for G . Clearly, $|M_G| \geq s$. Observe that for each removed vertex $a \in A_{\text{free}} \setminus (K \cup A_{\text{free}}^k)$ it holds that every vertex $a' \in A_{\text{free}}^k$ is to the right of a , that is, $a < a'$ and thus $N_{G-X}(a) \subseteq N_{G-X}(a')$. Since (G, s) is reduced with respect to Reduction Rule 4, it follows that $|M_{G-X}| \geq |M_G| - k$. Thus, there exist at most k augmenting paths for M_{G-X} in G . If none of these augmenting paths ends in a vertex $a \in A_{\text{free}} \setminus (K \cup A_{\text{free}}^k)$, then all augmenting paths exist also in G' and thus (G', s) is a yes-instance. If one of these augmenting paths, say P , ends in a , then at least one vertex $a' \in A_{\text{free}}^k$ is not endpoint of any of these augmenting paths. Since $a \notin K$, it follows from the definition of K that the neighbor b of a on P is indeed in $B \setminus X$. Since $N_{G-X}(a) \subseteq N_{G-X}(a')$, it follows that $\{a', b\} \in E$ and thus we can replace a by a' in the augmenting path. By exhaustively applying the above exchange argument, it follows that we can assume that none of the augmenting paths uses a vertex in $A_{\text{free}} \setminus (K \cup A_{\text{free}}^k)$. Thus, all augmenting paths are also contained in G' and hence the resulting instance (G', s) is still a yes-instance.

Finally observe that if (G', s) is a yes-instance, then also (G, s) is a yes-instance: any matching of size s in G' is also a matching in G since G' is a subgraph of G . \square

We have now all statements that we need to show our second main result.




Fig. 8 Schematic representation of the situation in the proof of Theorem 3; only the chain graph $G - X$ is shown. The vertices within the dashed boxes are bounded by the applications of Reduction Rules 5 to 6. Moreover, the vertices a_ℓ^K , a_r^K , b_ℓ^K , and b_r^K are all adjacent to vertices in X . It remains to upper-bound the vertices right of a_r^K and left of b_ℓ^K

Theorem 3 *MATCHING on bipartite graphs admits a cubic-vertex kernel with respect to the vertex deletion distance to chain graphs. The kernel can be computed in linear time.*

Proof Let (G, s) be the input instance with $G = (V, E)$, the two partitions $V = A \cup B$, and $X \subseteq V$ such that $G - X$ is a chain graph. If X is not given explicitly, then use the linear-time factor-four approximation provided in Lemma 13 to compute X . The kernelization is as follows: First, compute the matching M_{G-X} in linear time with Algorithm 4. Next compute the set of kept vertices K . Then, apply Reduction Rules 5 to 7. By Lemmas 17 and 18, this can be done in linear time. Let b_ℓ^K the leftmost vertex in $K \cap B$ and a_r^K rightmost vertex in $A \cap K$. Let a_ℓ^K and b_r^K be their matched neighbors. Since we reduced the instance with respect to Reduction Rule 5, we have $|K| \leq 2k^2$. Moreover, as we reduced the instance with respect to Reduction Rule 6, it follows that the number of vertices between a_ℓ^K and a_r^K as well as the number of vertices between b_ℓ^K and b_r^K is at most $4k^3$, respectively. Furthermore, there are at most $2k$ free vertices left in $V \setminus X$ since we reduced the instance with respect to Reduction Rule 7. It remains to upper-bound the number of matched vertices left of b_ℓ^K and right of a_r^K (see Fig. 8).

Observe that all vertices left of b_ℓ^K are matched with respect to M_{G-X} . If there are more than $2k$ vertices to the left of b_ℓ^K , then do the following: Add four vertices $a_\ell, b_\ell, x_\ell^a, x_\ell^b$ to V . The idea is that $\{a_\ell, b_\ell\}$ should be an edge in M_{G-X} such that $a_\ell \in A$ and $b_\ell \in B$ are in K and there is no vertex left of b_ℓ . This means we add these vertices to simulate the situation where the leftmost vertex in $B \setminus X$ is also in K . To ensure that a_ℓ and b_ℓ are in K and that they are not matched with some vertices in G , we add x_ℓ^a and x_ℓ^b to X and make x_ℓ^a respectively x_ℓ^b to their sole neighbors. In this way, we ensure that there is maximum matching in the new graph that is exactly two edges larger than the maximum matching in the old graph. In this new graph we can then apply Reduction Rule 6 to reduce the number of vertices between b_ℓ and b_ℓ^K . Formally, we add the following edges. Add $\{a_\ell, x_\ell^a\}, \{b_\ell, x_\ell^b\}$ to E . Add all edges between b_ℓ and the vertices in $B \setminus X$. Let a be the rightmost vertex in $A \setminus X$.

Then, add edges between a_ℓ and $N_{G-X}(a)$. Set $b' < b_\ell$ for each $b' \in B \setminus X$, set $a < a_\ell$ for each vertex $a \in A_{\text{free}}$, and set $a_\ell < a'$ for each matched vertex in $A \setminus X$. Furthermore, add $\{a_\ell, b_\ell\}$ to M_{G-X} and add a_ℓ and b_ℓ to K . Finally, increase s by two. Next, apply Reduction Rule 6 in linear time, then remove $a_\ell, b_\ell, x_\ell^a, x_\ell^b$ and reduce s by two. After this procedure, it follows that there are at most $2k$ vertices left of b_ℓ^K . If there are more than $2k$ vertices right of the rightmost vertex a_r^K in $A \cap K$, then use the same procedure as above. Thus, the total number of vertices in the remaining graph is at most $|X| + 2k + 4k^3 = O(k^3)$. Furthermore, observe that adding and removing the four vertices as well as applying Reduction Rule 6 can be done in linear time. Thus, the overall running time of the kernelization is $O(n + m)$. \square

Applying an $O(n^{2.5})$ -time algorithm for BIPARTITE MATCHING [15] on the kernel yields the following.

Corollary 3 *BIPARTITE MATCHING can be solved in $O(k^{7.5} + n + m)$ time, where k is the vertex deletion distance to chain graphs.*

Using the randomized $O(n^\omega)$ -time algorithm for MATCHING [22], one would obtain a randomized algorithm with running time $O(k^{3\omega} + n + m)$. Here, $\omega < 2.373$ is the matrix multiplication coefficient, that is, two $n \times n$ matrices can be multiplied in $O(n^\omega)$ time.

5 Conclusion

We focussed on kernelization results for unweighted (BIPARTITE) MATCHING. There remain numerous challenges for future research as discussed in the second part of this concluding section. First, however, let us discuss the closely connected issue of FPTP algorithms for MATCHING. There is a generic augmenting path-based approach to provide FPTP algorithms for MATCHING: Note that one can find an augmenting path in linear time [3, 11, 21]. Now the solving algorithm for MATCHING parameterized by some vertex deletion distance k works as follows:

1. Use a constant-factor linear-time (approximation) algorithm to compute a vertex set X such that $G - X$ is a “trivial” graph (where MATCHING is linear-time solvable).
2. Compute in linear time an initial maximum matching M in $G - X$.
3. Start with M as a matching in G and increase the size at most $|X| = k$ times to obtain in $O(k \cdot (n + m))$ time a maximum matching for G .

From this we can directly derive that MATCHING can be solved in $O(k(n + m))$ time, where k is one of the following parameters: feedback vertex number, feedback edge number, and vertex cover number. Moreover, BIPARTITE MATCHING can be solved in $O(k(n + m))$ time, where k is the vertex deletion distance to chain graphs. Using our kernelization results, the multiplicative dependence of the running time on

parameter k can now be made an additive one. For instance, in this way the running time for BIPARTITE MATCHING parameterized by vertex deletion distance to chain graphs “improves” from $O(k(n + m))$ to $O(k^{7.5} + n + m)$.

We conclude with some questions and tasks for future research. Can the size or the running time of the kernel with respect to feedback vertex set (see Sect. 3) be improved? In particular, can the exponential upper bound on the kernel size be decreased to a polynomial upper bound? Is there a linear-time computable kernel for MATCHING parameterized by the treedepth ℓ (assuming that ℓ is given)? This would complement the recent $O(\ell m)$ time algorithm [16]. Can one extend the kernel of Sect. 4 from BIPARTITE MATCHING to MATCHING parameterized by the distance to chain graphs?

Acknowledgements Open Access funding provided by Projekt DEAL. We thank anonymous reviewers of *Algorithmica* for their valuable feedback.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Bar-Yehuda, R., Geiger, D., Naor, J., Roth, R.M.: Approximation algorithms for the feedback vertex set problem with applications to constraint satisfaction and Bayesian inference. *SIAM J. Comput.* **27**(4), 942–959 (1998)
2. Bartha, M., Kresz, M.: A depth-first algorithm to reduce graphs in linear time. In: Proceedings of the 11th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC’09). IEEE, pp. 273–281 (2009)
3. Blum, N.: A new approach to maximum matching in general graphs. In: Proceedings of the 17th International Colloquium on Automata, Languages, and Programming (ICALP ’90). Vol. 443. LNCS. Springer, pp. 586–597 (1990)
4. Brandstädt, A., Le, V.B., Spinrad, J.P.: Graph Classes: a Survey, vol. 3. SIAM, SIAM Monographs on Discrete Mathematics and Applications (1999)
5. Cai, L.: Parameterized complexity of vertex colouring. *Discrete Appl. Math.* **127**(1), 415–429 (2003)
6. Chang, M.: Algorithms for maximum matching and minimum fill-in on Chordal bipartite graphs. In: Proceedings of the 7th International Symposium on Algorithms and Computation (ISAAC ’96). Vol. 1178. LNCS. Springer, pp. 146–155 (1996)
7. Dahlhaus, E., Karpinski, M.: Matching and multidimensional matching in chordal and strongly chordal graphs. *Discrete Appl. Math.* **84**(1–3), 79–91 (1998)
8. Duan, R., Pettie, S.: Linear-time approximation for maximum weight matching. *J. ACM* **61**(1), 1:1–1:23 (2014)
9. Fomin, F.V., Lokshtanov, D., Pilipeczuk, M., Saurabh, S., Wrochna, M.: Fully polynomial-time parameterized computations for graphs and matrices of low Treewidth. *ACM Trans. Algorithms* **14**(3), 34:1–34:45 (2018)

10. Gabow, H.N., Tarjan, R.E.: A linear-time algorithm for a special case of disjoint set union. *J. Comput. Syst. Sci.* **30**(2), 209–221 (1985)
11. Gabow, H.N., Tarjan, R.E.: Faster scaling algorithms for general graph-matching problems. *J. ACM* **38**(4), 815–853 (1991)
12. Giannopoulou, A.C., Mertzios, G.B., Niedermeier, R.: Polynomial fixed-parameter algorithms: a case study for longest path on interval graphs. *Theor. Comput. Sci.* **689**, 67–95 (2017)
13. Guo, J., Hüffner, F., Niedermeier, R.: A structural view on parameterizing problems: distance from triviality. In: *Proceedings of the 1st International Workshop on Parameterized and Exact Computation (IWPEC '04)*. Vol. 3162. LNCS. Springer, pp. 162–173 (2004)
14. Gupta, M., Peng, R.: Fully Dynamic $(1 + \epsilon)$ -Approximate Matchings. In: *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS '13)*. IEEE, pp. 548–557 (2013)
15. Hopcroft, J.E., Karp, R.M.: An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM J. Comput.* **2**(4), 225–231 (1973)
16. Iwata, Y., Ogasawara, T., Ohsaka, N.: On the power of tree-depth for fully polynomial FPT algorithms. In: *35th Symposium on Theoretical Aspects of Computer Science (STACS '18)*. Vol. 96. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 41:1–41:14 (2018)
17. Karp, R. M., Sipser, M.: Maximum matchings in sparse random graphs. In: *Proceedings of the 22nd Annual IEEE Symposium on Foundations of Computer Science (FOCS '81)*. IEEE, pp. 364–375 (1981)
18. Korenwein, V., Nichterlein, A., Niedermeier, R., Zschoche, P.: Data reduction for maximum matching on real-world graphs: theory and experiments. In: *Proceedings of the 26th Annual European Symposium on Algorithms (ESA 2018)*. Vol. 112. LIPIcs. Schloss Dagstuhl - Leibniz- Zentrum fuer Informatik, 53:1–53:13 (2018)
19. Kratsch, S., Nelles, F.: Efficient and Adaptive Parameterized Algorithms on Modular Decompositions. In: *Proceedings of the 26th Annual European Symposium on Algorithms (ESA 2018)*. Vol. 112. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 55:1–55:15 (2018)
20. Mertzios, G.B., Nichterlein, A., Niedermeier, R.: Linear-time algorithm for maximum-cardinality matching on cocomparability graphs. *SIAM J. Discrete Math.* **32**(4), 2820–2835 (2018)
21. Micali, S., Vazirani, V. V.: An $O(\sqrt{|V||E|})$ algorithm for finding maximum matching in general graphs. In: *Proceedings of the 21st Annual IEEE Symposium on Foundations of Computer Science (FOCS '80)*. IEEE, pp. 17–27 (1980)
22. Mucha, M., Sankowski, P.: Maximum matchings via Gaussian elimination. In: *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS '04)*. IEEE, pp. 248–255 (2004)
23. Nešetřil, J., de Mendez, P.O.: *Sparsity—Graphs, Structures, and Algorithms*, vol. 28. Springer, Algorithms and Combinatorics (2012)
24. Skiena, S.S.: *The Algorithm Design Manual*. Springer, Berlin (2010)
25. Steiner, G., Yeomans, J.S.: A linear time algorithm for maximum matchings in convex bipartite graphs. *Comput. Math. Appl.* **31**, 91–96 (1996)
26. Yuster, R.: Maximum matching in regular and almost regular graphs. *Algorithmica* **66**(1), 87–92 (2013)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.