# Deterministic Fully Dynamic Data Structures for Vertex Cover and Matching* 

Sayan Bhattacharya ${ }^{\dagger} \quad$ Monika Henzinger ${ }^{\ddagger} \quad$ Giuseppe F. Italiano ${ }^{\text {§ }}$

December 4, 2014


#### Abstract

We present the first deterministic data structures for maintaining approximate minimum vertex cover and maximum matching in a fully dynamic graph $G=(V, E)$, with $|V|=n$ and $|E|=m$, in $o(\sqrt{m})$ time per update. In particular, for minimum vertex cover we provide deterministic data structures for maintaining a $(2+\epsilon)$ approximation in $O\left(\log n / \epsilon^{2}\right)$ amortized time per update. For maximum matching, we show how to maintain a $(3+\epsilon)$ approximation in $O\left(\min \left(\sqrt{n} / \epsilon, m^{1 / 3} / \epsilon^{2}\right)\right)$ amortized time per update, and a $(4+\epsilon)$ approximation in $O\left(m^{1 / 3} / \epsilon^{2}\right)$ worst-case time per update. Our data structure for fully dynamic minimum vertex cover is essentially near-optimal and settles an open problem by Onak and Rubinfeld [13].


[^0]
## Contents

1 Introduction ..... 1
1.1 Previous work. ..... 1
1.2 Our results ..... 2
1.3 Our techniques ..... 2
2 Deterministic Fully Dynamic Vertex Cover ..... 3
2.1 The ( $\alpha, \beta$ )-partition and its properties. ..... 4
2.2 Handling the insertion/deletion of an edge ..... 6
2.3 Data structures. ..... 7
2.4 Bounding the amortized update time: An overview ..... 7
2.5 Implementation details ..... 9
2.6 Bounding the amortized update time: Detailed analysis ..... 14
2.6.1 Analysis of the FIX $(v)$ subroutine ..... 15
3 Dynamic Matching: Preliminaries ..... 19
3.1 Maintaining an approximate matching in a bounded degree graph. ..... 20
3.1.1 Proof of Theorem 3.4 ..... 21
3.2 The kernel and its properties. ..... 23
3.2.1 Proof of Theorem 3.9. ..... 23
3.2.2 Proof of Theorem 3.10. ..... 24
3.2.3 Proof of Theorem 3.11 ..... 25
3.3 An algorithm for building a kernel in a static graph. ..... 26
3.4 Data structures for representing a kernel in a dynamic graph. ..... 26
$4(3+\epsilon)$-approximate matching in $O(\sqrt{n} / \epsilon)$ amortized update time ..... 27
4.1 Overview of our approach. ..... 27
4.2 Maintaining the kernel $\kappa(G)$. ..... 27
4.3 Maintaining the matching $M$ in $\kappa(G)$. ..... 29
$5(3+\epsilon)$-approximate matching in $O\left(m^{1 / 3} / \epsilon^{2}\right)$ amortized update time ..... 30
5.1 Overview of our approach. ..... 30
5.2 Algorithms for maintaining the kernel during a phase. ..... 31
5.2.1 Handling an edge insertion in $G$ during the phase. ..... 31
5.2.2 Handling an edge deletion in $G$ during the phase. ..... 32
5.2.3 Proof of Theorem 5.2. ..... 32
$6(4+\epsilon)$-approximate matching in $O\left(m^{1 / 3} / \epsilon^{2}\right)$ worst-case update time ..... 34
6.1 Overview of our approach. ..... 34
6.2 Algorithm for maintaining and building the kernel during each phase. ..... 35
6.2.1 Dealing with the current phase ..... 35
6.2.2 Preparing for the next phase. ..... 36
6.2.3 Handling an edge insertion into $G$ during the epoch. ..... 37
6.2.4 Handling an edge deletion in $G$ during the epoch. ..... 37
6.2.5 Proof of Theorem 6.3 ..... 38

## 1 Introduction

Finding maximum matchings and minimum vertex covers in undirected graphs are classical problems in combinatorial optimization. Let $G=(V, E)$ be an undirected graph, with $m=|E|$ edges and $n=|V|$ nodes. A matching in $G$ is a set of vertex-disjoint edges, i.e., no two edges share a common vertex. A maximum matching, also known as maximum cardinality matching, is a matching with the largest possible number of edges. A matching is maximal if it is not a proper subset of any other matching in $G$. A subset $V^{\prime} \subseteq V$ is a vertex cover if each edge of $G$ is incident to at least one vertex in $V^{\prime}$. A minimum vertex cover is a vertex cover of smallest possible size.

The Micali-Vazirani algorithm for maximum matching runs in $O(m \sqrt{n})$ time [5, 10]. Using this algorithm, a $(1+\epsilon)$-approximate maximum matching can be constructed in $O(m / \epsilon)$ time [3]. Finding a minimum vertex cover, on the other hand, is NP-hard. Still, these two problems remain closely related as their LP-relaxations are duals of each other. Furthermore, a maximal matching, which can be computed in $O(m)$ time in a greedy fashion, is known to provide a 2-approximation both to maximum matching and to minimum vertex cover (by using the endpoints of the maximal matching). Under the unique games conjecture, the minimum vertex cover cannot be efficiently approximated within any constant factor better than 2 [7]. Thus, under the unique games conjecture, the 2-approximation in $O(m)$ time by the greedy method is the optimal guarantee for this problem.

In this paper, we consider a dynamic setting, where the input graph is being updated via a sequence of edge insertions/deletions. The goal is to design data structures that are capable of maintaining the solution to an optimization problem faster than recomputing it from scratch after each update. If $P \neq N P$ we cannot achieve polynomial time updates for minimum vertex cover. We also observe that achieving fast update times for maximum matching appears to be a particularly difficult task: in this case, an update bound of $O$ (poly $\log (n)$ ) would be a breakthrough, since it would immediately improve the longstanding bounds of various static algorithms [5, 9-11]. The best known update bound for dynamic maximum matching is obtained by a randomized data structure of Sankowski [14], which has $O\left(n^{1.495}\right)$ time per update. In this scenario, if one wishes to achieve fast update times for dynamic maximum matching or minimum vertex cover, approximation appears to be inevitable. Indeed, in the last few years there has been a growing interest in designing efficient dynamic data structures for maintaining approximate solutions to both these problems.

### 1.1 Previous work.

A maximal matching can be maintained in $O(n)$ worst-case update time by a trivial deterministic algorithm. Ivković and Lloyd [6] showed how to improve this bound to $O\left((n+m)^{\sqrt{2} / 2}\right)$. Onak and Rubinfeld [13] designed a randomized data structure that maintains constant factor approximations to maximum matching and to minimum vertex cover in $O\left(\log ^{2} n\right)$ amortized time per update with high probability, with the approximation factors being large constants. Baswana, Gupta and Sen [2] improved these bounds by showing that a maximal matching, and thus a 2 -approximation of maximum matching and minimum vertex cover, can be maintained in a dynamic graph in amortized $O(\log n)$ update time with high probability.

Subsequently, turning to deterministic data structures, Neiman and Solomon [12] showed that a $3 / 2$ approximate maximum matching can be maintained dynamically in $O(\sqrt{m})$ worst-case time per update. Their data structure maintains a maximal matching and thus achieves the same update bound also for 2approximate minimum vertex cover. Furthermore, Gupta and Peng [4] presented a deterministic data structure to maintain a $(1+\epsilon)$ approximation of a maximum matching in $O\left(\sqrt{m} / \epsilon^{2}\right)$ worst-case time per update. We also note that Onak and Rubinfeld [13] gave a deterministic data structure that maintains an $O(\log n)$ approximate minimum vertex cover in $O\left(\log ^{2} n\right)$ amortized update time. See Table 1 for a summary of these results.

Very recently, Abboud and Vassilevska Williams [1] showed a conditional lower bound on the perfor-
mance of any dynamic matching algorithm. There exists an integer $k \in[2,10]$ with the following property: if the dynamic algorithm maintains a matching with the property that every augmenting path in the input graph (w.r.t. the matching) has length at least $(2 k-1)$, then an amortized update time of $o\left(m^{1 / 3}\right)$ for the algorithm will violate the 3 -SUM conjecture (which states that the 3 -SUM problem on $n$ numbers cannot be solved in $o\left(n^{2}\right)$ time $)$.

### 1.2 Our results.

From the above discussion, it is clear that for both fully dynamic constant approximate maximum matching and minimum vertex cover, there is a huge gap between state of the art deterministic and randomized performance guarantees: the former gives $O(\sqrt{m})$ update time, while the latter gives $O(\log n)$ update time. Thus, it seems natural to ask whether the $O(\sqrt{m})$ bound achieved in $[4,12]$ is a natural barrier for deterministic data structures. In particular, in their pioneering work on these problems, Onak and Rubinfeld [13] asked:

- "Is there a deterministic data structure that achieves a constant approximation factor with polylogarithmic update time?"

We answer this question in the affirmative by presenting a deterministic data structure that maintains a $(2+\epsilon)$-approximation of a minimum vertex cover in $O\left(\log n / \epsilon^{2}\right)$ amortized time per update. Since it is impossible to get better than 2 -approximation for minimum vertex cover in polynomial time, our data structure is near-optimal (under the unique games conjecture). As a by product of our approach, we can also maintain, deterministically, a $(2+\epsilon)$-approximate maximum fractional matching in $O\left(\log n / \epsilon^{2}\right)$ amortized update time. Note that the vertices of the fractional matching polytope of a graph are known to be half integral, i.e., they have only $\{0,1 / 2,1\}$ coordinates (see, e.g., [8]). This implies immediately that the value of any fractional matching is at most $3 / 2$ times the value of the maximum integral matching. Thus, it follows that we can maintain the value of the maximum (integral) matching within a factor of $(2+\epsilon) \cdot(3 / 2)=$ $(3+O(\epsilon))$, deterministically, in $O\left(\log n / \epsilon^{2}\right)$ amortized update time.

Next, we focus on the problem of maintaining an integral matching in a dynamic setting. For this problem, we show how to maintain a $(3+\epsilon)$-approximate maximum matching in $O\left(\min \left(\sqrt{n} / \epsilon, m^{1 / 3} / \epsilon^{2}\right)\right)$ amortized time per update, and a $(4+\epsilon)$-approximate maximum matching in $O\left(m^{1 / 3} / \epsilon^{2}\right)$ worst-case time per update. Since $m^{1 / 3}=o(n)$, we provide the first deterministic data structures for dynamic matching whose update time is sublinear in the number of nodes. Table 1 puts our main results in perspective with previous work.

### 1.3 Our techniques.

To see why it is difficult to deterministically maintain a dynamic (say maximal) matching, consider the scenario when a matched edge incident to a node $u$ gets deleted from the graph. To recover from this deletion, we have to scan through the adjacency list of $u$ to check if it has any free neighbor $z$. This takes time proportional to the degree of $u$, which can be $O(n)$. Both the papers [2, 13] use randomization to circumvent this problem. Roughly speaking, the idea is to match the node $u$ to one of its free neighbors $z$ picked at random, and show that even if this step takes $O(\operatorname{deg}(u))$ time, in expectation the newly matched edge $(u, z)$ survives the next $\operatorname{deg}(u) / 2$ edge deletions in the graph (assuming that the adversary is not aware of the random choices made by the data structure). This is used to bound the amortized update time.

Our key insight is that we can maintain a large fractional matching deterministically. Suppose that in this fractional matching, we pick each edge incident to $u$ to an extent of $1 / \operatorname{deg}(u)$. These edges together contribute at most one to the objective. Thus, we do not have to do anything for the next deg $(u) / 2$ edge deletions incident to $u$, as these deletions reduce the contribution of $u$ towards the objective by at most a factor of two. This gives us the desired amortized bound. Inspired by this observation, we take a closer look

| Problem | Approximation <br> Guarantee | Update <br> Time | Data <br> Structure | Reference |
| :---: | :---: | :---: | :---: | :---: |
| MM \& MVC | $O(1)$ | $O\left(\log ^{2} n\right)$ amortized | randomized | $[13]$ |
| MM \& MVC | 2 | $O(\log n)$ amortized | randomized | $[2]$ |
| MM | 1.5 | $O(\sqrt{m})$ worst-case | deterministic | $[12]$ |
| MVC | 2 | $O(\sqrt{m})$ worst-case | deterministic | $[12]$ |
| MM | $1+\epsilon$ | $O\left(\sqrt{m} / \epsilon^{2}\right)$ worst-case | deterministic | $[4]$ |
| MVC | $2+\epsilon$ | $O\left(\log n / \epsilon^{2}\right)$ amortized | deterministic | This paper |
| MM | $3+\epsilon$ | $O(\sqrt{n} / \epsilon)$ amortized | deterministic | This paper |
| MM | $3+\epsilon$ | $O\left(m^{1 / 3} / \epsilon^{2}\right)$ amortized | deterministic | This paper |
| MM | $4+\epsilon$ | $O\left(m^{1 / 3} / \epsilon^{2}\right)$ worst-case | deterministic | This paper |

Table 1: Dynamic data structures for approximate (integral) maximum matching (MM) and minimum vertex cover (MVC).
at the framework of Onak and Rubinfeld [13]. Roughly speaking, they maintain a hierarchical partition of the set of nodes $V$ into $O(\log n)$ levels such that the nodes in all but the lowest level, taken together, form a valid vertex cover $V^{*}$. In addition, they maintain a matching $M^{*}$ as a dual certificate. Specifically, they show that $\left|V^{*}\right| \leq \lambda \cdot\left|M^{*}\right|$ for some constant $\lambda$, which implies that $V^{*}$ is a $\lambda$-approximate minimum vertex cover. Their data structure is randomized since, as discussed above, it is particularly difficult to maintain the matching $M^{*}$ deterministically in a dynamic setting. To make the data structure deterministic, instead of $M^{*}$, we maintain a fractional matching as a dual certificate. Along the way, we improve the amortized update time of [13] from $O\left(\log ^{2} n\right)$ to $O\left(\log n / \epsilon^{2}\right)$, and their approximation guarantee from some large constant $\lambda$ to $2+\epsilon$.

Our approach gives near-optimal bounds for fully dynamic minimum vertex cover and, as we have already remarked, it maintains a fractional matching. Next, we consider the problem of maintaining an approximate maximum integral matching, for which we are able to provide deterministic data structures with improved (polynomial) update time. Towards this end, we introduce the concept of a kernel of a graph, which we believe is of independent interest. Intuitively, a kernel is a subgraph with two important properties: (i) each node has bounded degree in the kernel, and (ii) a kernel approximately preserves the size of the maximum matching in the original graph. Our key contribution is to show that a kernel always exists, and that it can be maintained efficiently in a dynamic graph undergoing a sequence of edge updates.

## 2 Deterministic Fully Dynamic Vertex Cover

The input graph $G=(V, E)$ has $|V|=n$ nodes and zero edges in the beginning. Subsequently, it keeps getting updated due to the insertions of new edges and the deletions of already existing edges. The edge updates, however, occur one at a time, while the set $V$ remains fixed. The goal is to maintain an approximate vertex cover of $G$ in this fully dynamic setting.

In Section 2.1, we introduce the notion of an $(\alpha, \beta)$-partition of $G=(V, E)$. This is a hierarchical partition of the set $V$ into $L+1$ levels, where $L=\left\lceil\log _{\beta}(n / \alpha)\right\rceil$ and $\alpha, \beta>1$ are two parameters (Definition 2.2). If the ( $\alpha, \beta$ )-partition satisfies an additional property (Invariant 2.4), then from it we can easily derive a $2 \alpha \beta$-approximation to the minimum vertex cover (Theorem 2.5).

In Section 2.2, we outline a natural deterministic algorithm for maintaining such an $(\alpha, \beta)$-partition. In Section 2.3, we present the relevant data structures that are required for implementing the algorithm. We
will analyze its amortized update time using an extremely fine tuned potential function: In Section 2.4, we present a high level overview of this approach. Section 2.5 contains a detailed implementation of the algorithm using the data structures from Section 2.3. Finally, in Section 2.6 , we give a complete analysis of the amortized update time of the algorithm. Specifically, we show that for $\alpha=1+3 \epsilon$ and $\beta=1+\epsilon$, the algorithm takes $O\left((t / \epsilon) \log _{1+\epsilon} n\right)$ time to handle $t$ edge updates starting from an empty graph (see Theorem 2.11). This leads to the main result of this section, which is summarized in the theorem below.

Theorem 2.1. For every $\epsilon \in(0,1]$, we can deterministically maintain a $(2+\epsilon)$-approximate vertex cover in a fully dynamic graph, the amortized update time being $O\left(\log n / \epsilon^{2}\right)$.

### 2.1 The $(\alpha, \beta)$-partition and its properties.

Definition 2.2. An $(\alpha, \beta)$-partition of the graph $G$ partitions its node-set $V$ into subsets $V_{0} \ldots V_{L}$, where $L=\left\lceil\log _{\beta}(n / \alpha)\right\rceil$ and $\alpha, \beta>1$. For $i \in\{0, \ldots, L\}$, we identify the subset $V_{i}$ as the $i^{\text {th }}$ level of this partition, and denote the level of a node $v$ by $\ell(v)$. Thus, we have $v \in V_{\ell(v)}$ for all $v \in V$. Furthermore, the partition assigns a weight $w(u, v)=\beta^{-\max (\ell(u), \ell(v))}$ to every edge $(u, v) \in V$.

Define $\mathcal{N}_{v}$ to be the set of neighbors of a node $v \in V$. Given an $(\alpha, \beta)$-partition, let $\mathcal{N}_{v}(i) \subseteq \mathcal{N}_{v}$ denote the set of neighbors of $v$ that are in the $i^{\text {th }}$ level, and let $\mathcal{N}_{v}(i, j) \subseteq \mathcal{N}_{v}$ denote the set of neighbors of $v$ whose levels are in the range $[i, j]$.

$$
\begin{gather*}
\mathcal{N}_{v}=\{u \in V:(u, v) \in E\} \forall v \in V .  \tag{1}\\
\mathcal{N}_{v}(i)=\left\{u \in \mathcal{N}_{v} \cap V_{i}\right\} \forall v \in V ; i \in\{0, \ldots, L\}  \tag{2}\\
\mathcal{N}_{v}(i, j)=\bigcup_{k=i}^{j} \mathcal{N}_{v}(k) \forall v \in V ; i, j \in\{0, \ldots, L\}, i \leq j . \tag{3}
\end{gather*}
$$

Similarly, define the notations $D_{v}, D_{v}(i)$ and $D_{v}(i, j)$. Note that $D_{v}$ is the degree of a node $v \in V$.

$$
\begin{align*}
D_{v} & =\left|\mathcal{N}_{v}\right|  \tag{4}\\
D_{v}(i) & =\left|\mathcal{N}_{v}(i)\right|  \tag{5}\\
D_{v}(i, j) & =\left|\mathcal{N}_{v}(i, j)\right| \tag{6}
\end{align*}
$$

Given an $(\alpha, \beta)$-partition, let $W_{v}=\sum_{u \in \mathcal{N}_{v}} w(u, v)$ denote the total weight a node $v \in V$ receives from the edges incident to it. We also define the notation $W_{v}(i)$. It gives the total weight the node $v$ would receive from the edges incident to it, if the node $v$ itself were to go to the $i^{\text {th }}$ level. Thus, we have $W_{v}=W_{v}(\ell(v))$. Since the weight of an edge $(u, v)$ in the hierarchical partition is given by $w(u, v)=\beta^{-\max (\ell(u), \ell(v))}$, we derive the following equations for all nodes $v \in V$.

$$
\begin{gather*}
W_{v}=\sum_{u \in \mathcal{N}_{v}} \beta^{-\max (\ell(u), \ell(v))} .  \tag{7}\\
W_{v}(i)=\sum_{u \in \mathcal{N}_{v}} \beta^{-\max (\ell(u), i)} \forall i \in\{0, \ldots, L\} . \tag{8}
\end{gather*}
$$

Lemma 2.3. Every $(\alpha, \beta)$-partition of the graph $G$ satisfies the following conditions for all nodes $v \in V$.

$$
\begin{gather*}
W_{v}(L) \leq \alpha  \tag{9}\\
W_{v}(L) \leq \cdots \leq W_{v}(i) \leq \cdots \leq W_{v}(0)  \tag{10}\\
W_{v}(i) \leq \beta \cdot W_{v}(i+1) \forall i \in\{0, \ldots, L-1\} . \tag{11}
\end{gather*}
$$

Proof. Fix any $(\alpha, \beta)$-partition and any node $v \in V$. We prove the first part of the lemma as follows.

$$
\begin{array}{r}
W_{v}(L)=\sum_{u \in \mathcal{N}_{v}} \beta^{-\max (\ell(u), L)} \\
=\sum_{u \in \mathcal{N}_{v}} \beta^{-L} \leq n \cdot \beta^{-L} \leq n \cdot \beta^{-\log _{\beta}(n / \alpha)}=\alpha .
\end{array}
$$

We now fix any level $i \in\{0, \ldots, L-1\}$ and show that the $(\alpha, \beta)$-partition satisfies equation 10 .

$$
\begin{aligned}
& W_{v}(i+1)=\sum_{u \in \mathcal{N}_{v}} \beta^{-\max (\ell(u), i+1)} \\
& \quad \leq \sum_{u \in \mathcal{N}_{v}} \beta^{-\max (\ell(u), i)}=W_{v}(i)
\end{aligned}
$$

Finally, we prove equation 11.

$$
\begin{aligned}
W_{v}(i)= & \sum_{u \in \mathcal{N}_{v}} \beta^{-\max (\ell(u), i)}=\beta \cdot \sum_{u \in \mathcal{N}_{v}} \beta^{-1-\max (\ell(u), i)} \\
& \leq \beta \cdot \sum_{u \in \mathcal{N}_{v}} \beta^{-\max (\ell(u), i+1)}=\beta \cdot W_{v}(i+1)
\end{aligned}
$$

Fix any node $v \in V$, and focus on the value of $W_{v}(i)$ as we go down from the highest level $i=L$ to the lowest level $i=0$. Lemma 2.3 states that $W_{v}(i) \leq \alpha$ when $i=L$, that $W_{v}(i)$ keeps increasing as we go down the levels one after another, and that $W_{v}(i)$ increases by at most a factor of $\beta$ between consecutive levels.

We will maintain a specific type of $(\alpha, \beta)$-partition, where each node is assigned to a level in a way that satisfies Invariant 2.4.

Invariant 2.4. For every node $v \in V$, if $\ell(v)=0$, then $W_{v} \leq \alpha \cdot \beta$. Else if $\ell(v) \geq 1$, then $W_{v} \in[1, \alpha \beta]$.
Consider any $(\alpha, \beta)$-partition satisfying Invariant 2.4. Let $v \in V$ be a node in this partition that is at level $\ell(v)=k \in\{0, \ldots, L\}$. It follows that $\sum_{u \in \mathcal{N}_{v}(0, k)} w(u, v)=\left|\mathcal{N}_{v}(0, k)\right| \cdot \beta^{-k} \leq W_{v} \leq \alpha \beta$. Thus, we infer that $\left|\mathcal{N}_{v}(0, k)\right| \leq \alpha \beta^{k+1}$. In other words, Invariant 2.4 gives an upper bound on the number of neighbors a node $v$ can have that lie on or below $\ell(v)$. We will crucially use this property in the analysis of our algorithm.

Theorem 2.5. Consider an $(\alpha, \beta)$-partition of the graph $G$ that satisfies Invariant 2.4. Let $V^{*}=\{v \in V$ : $\left.W_{v} \geq 1\right\}$ be the set of nodes with weight at least one. The set $V^{*}$ is a feasible vertex cover in $G$. Further, the size of the set $V^{*}$ is at most $2 \alpha \beta$-times the size of the minimum-cardinality vertex cover in $G$.

Proof. Consider any edge $(u, v) \in E$. We claim that at least one of its endpoints belong to the set $V^{*}$. Suppose that the claim is false and we have $W_{u}<1$ and $W_{v}<1$. If this is the case, then Invariant 2.4 implies that $\ell(u)=\ell(v)=0$ and $w(u, v)=\beta^{-\max (\ell(u), \ell(v))}=1$. Since $W_{u} \geq w(u, v)$ and $W_{v} \geq w(u, v)$, we get $W_{u} \geq 1$ and $W_{v} \geq 1$, and this leads to a contradiction. Thus, we infer that the set $V^{*}$ is a feasible vertex cover in the graph $G$.

Next, we construct a fractional matching $M_{f}$ by picking every edge $(u, v) \in E$ to an extent of $x(u, v)=$ $w(u, v) /(\alpha \beta) \in[0,1]$. Since for all nodes $v \in V$, we have $\sum_{u \in \mathcal{N}_{v}} x(u, v)=\sum_{u \in \mathcal{N}_{v}} w(u, v) /(\alpha \beta)=$
$W_{v} /(\alpha \beta) \leq 1$, we infer that $M_{f}$ is a valid fractional matching in $G$. The size of this matching is given by $\left|M_{f}\right|=\sum_{(u, v) \in E} x(u, v)=(1 /(\alpha \beta)) \cdot \sum_{(u, v) \in E} w(u, v)$. We now bound the size of $V^{*}$ in terms of $\left|M_{f}\right|$.

$$
\begin{array}{r}
\left|V^{*}\right|=\sum_{v \in V^{*}} 1 \leq \sum_{v \in V^{*}} W_{v}=\sum_{v \in V^{*}} \sum_{u \in \mathcal{N}_{v}} w(u, v) \\
\leq \sum_{v \in V} \sum_{u \in \mathcal{N}_{v}} w(u, v)=2 \cdot \sum_{(u, v) \in E} w(u, v)=(2 \alpha \beta) \cdot\left|M_{f}\right|
\end{array}
$$

The approximation guarantee now follows from the LP duality between minimum fractional vertex cover and maximum fractional matching.

Query time. We store the nodes $v$ with $W_{v} \geq 1$ as a separate list. Thus, we can report the set of nodes in the vertex cover in $O(1)$ time per node. Using appropriate pointers, we can report in $O(1)$ time whether or not a given node is part of this vertex cover. In $O(1)$ time we can also report the size of the vertex cover.

### 2.2 Handling the insertion/deletion of an edge.

A node is called dirty if it violates Invariant 2.4, and clean otherwise. Since the graph $G=(V, E)$ is initially empty, every node is clean and at level zero before the first update in $G$. Now consider the time instant just prior to the $t^{t h}$ update in $G$. By induction hypothesis, at this instant every node is clean. Then the $t^{t h}$ update takes place, which inserts (resp. deletes) an edge $(x, y)$ in $G$ with weight $w(x, y)=\beta^{-\max (\ell(x), \ell(y))}$. This increases (resp. decreases) the weights $W_{x}, W_{y}$ by $w(x, y)$. Due to this change, the nodes $x$ and $y$ might become dirty. To recover from this, we call the subroutine in Figure 1.

```
01. While there exists a dirty node \(v\)
02. IF \(W_{v}>\alpha \beta\), THEN
    // If true, then by equation \(9 \ell(v)<L\).
03. Increment the level of \(v\)
        by setting \(\ell(v) \leftarrow \ell(v)+1\).
04. ELSE IF \(\left(W_{v}<1\right.\) and \(\left.\ell(v)>0\right)\), THEN
05. Decrement the level of \(v\)
    by setting \(\ell(v) \leftarrow \ell(v)-1\).
```

Figure 1: RECOVER().

Consider any node $v \in V$ and suppose that $W_{v}=W_{v}(\ell(v))>\alpha \beta$. In this event, equation 9 implies that $W_{v}(L)<W_{v}(\ell(v))$ and hence we have $L>\ell(v)$. In other words, when the procedure described in Figure 1 decides to increment the level of a dirty node $v$ (Step 03), we know for sure that the current level of $v$ is strictly less than $L$ (the highest level in the $(\alpha, \beta)$-partition).

Next, consider a node $z \in \mathcal{N}_{v}$. If we change $\ell(v)$, then this may change the weight $w(v, z)$, and this in turn may change the weight $W_{z}$. Thus, a single iteration of the While loop in Figure 1 may lead to some clean nodes becoming dirty, and some other dirty nodes becoming clean. If and when the While loop terminates, however, we are guaranteed that every node is clean and that Invariant 2.4 holds.

Comparison with the framework of Onak and Rubinfeld [13]. As described below, there are two significant differences between our framework and that of [13]. Consequently, many of the technical details of our approach (illustrated in Section 2.6) differ from the proof in [13].

First, in the hierarchical partition of [13], the invariant for a node $y$ consists of $O(L)$ constraints: for each level $i \in\{\ell(y), \ldots, L\}$, the quantity $\left|\mathcal{N}_{y}(0, i)\right|$ has to lie within a certain range. This is the main reason
for their amortized update time being $\Theta\left(\log ^{2} n\right)$. Indeed when a node $y$ becomes dirty, unlike in our setting, they have to spend $\Theta(\log n)$ time just to figure out the new level of $y$.

Second, along with the hierarchical partition, the authors in [13] maintain a matching as a dual certificate, and show that the size of this matching is within a constant factor of the size of their vertex cover. As pointed out in Section 1, this is the part where they crucially need to use randomization, as till date there is no deterministic data structure for maintaining a large matching in polylog amortized update time. We bypass this barrier by implicitly maintaining a fractional matching as a dual certificate. Indeed, the weight $w(y, z)$ of an edge $(y, z)$ in our hierarchical partition, after suitable scaling, equals the fractional extent by which the edge $(y, z)$ is included in our fractional matching.

### 2.3 Data structures.

We now describe the data structures that we use to implement the algorithm outlined in Section 2.2

- We maintain the following data structures for each node $v \in V$.
- A counter Level $[v]$ to keep track of the current level of $v$. Thus, we set Level $[v] \leftarrow \ell(v)$.
- A counter Weight $[v]$ to keep track of the weight of $v$. Thus, we set $\operatorname{Weight}[v] \leftarrow W_{v}$.
- For every level $i>\operatorname{Level}[v]$, the set of nodes $\mathcal{N}_{v}(i)$ as a doubly linked list $\operatorname{Neighbors~}_{v}[i]$. For every level $i \leq \operatorname{Level}[v]$, the list Neighbors ${ }_{v}[i]$ is empty.
- For level $i=\operatorname{Level}[v]$, the set of nodes $\mathcal{N}_{v}(0, i)$ as a doubly linked list Neighbors ${ }_{v}[0, i]$. For every level $i \neq \operatorname{Level}[v]$, the list Neighbors ${ }_{v}[0, i]$ is empty.
- When the graph $G$ gets updated due to an edge insertion/deletion, we may discover that a node violates Invariant 2.4. Such a node is called dirty, and we store the set of such nodes as a doubly linked list Dirty-nodes. For every node $v \in V$, we maintain a bit Status $[v] \in\{$ dirty, clean $\}$ that indicates if the node is dirty or not. Every dirty node stores a pointer to its position in the list Dirty-nodes.
- The phrase "neighborhood lists of $v$ " refers to the set $\bigcup_{i=0}^{L}\left\{\right.$ NEIGHBors $\left._{v}[0, i], \operatorname{NEIGHBORS}_{v}[i]\right\}$. For every edge $(u, v)$, we maintain two bidirectional pointers: one links the edge to the position of $v$ in the neighborhood lists of $u$, while the other links the edge to the position of $u$ in the neighborhood lists of $v$.


### 2.4 Bounding the amortized update time: An overview

In Section 2.5, we present a detailed implementation of our algorithm using the data structures described in Section 2.3. In Section 2.6, we prove that for any $\epsilon \in[0,1], \alpha=1+3 \epsilon$ and $\beta=1+\epsilon$, it takes $O\left(t \log n / \epsilon^{2}\right)$ time to handle $t$ edge insertions/deletions in $G$ starting from an empty graph. This gives an amortized update time of $O\left(\log n / \epsilon^{2}\right)$, and by Theorem 2.5, a $(2+14 \epsilon)$-approximation to the minimum vertex cover in $G$. The proof works as follows. First, we note that after an edge insertion or deletion the data structure can be updated in time $O(1)$ plus the time to adjust the levels of the nodes, i.e., the time for procedure RECOVER (see Figure 1). To bound the latter we show that it takes time $\Theta\left(1+D_{v}(0, i)\right)$, when node $v$ changes from level $i$ to level $i+1$ or level $i-1$, and prove the bound on the total time spent in procedure RECOVER using a potential function based argument.

As the formal analysis is quite involved, some high level intuitions are in order. Accordingly, in this section, we describe the main idea behind a (slightly) simplified variant of the above argument, which gives an amortized bound on the number of times we have to change the weight of an already existing edge. This
number is $D_{v}(0, i)$, when node $v$ changes from level $i$ to level $i+1$ and $D_{v}(0, i-1)$, when node $v$ changes from level $i$ to level $i-1 .{ }^{1}$

To simplify the exposition even further, we assume that $\alpha, \beta$ are sufficiently large constants, and describe the main idea behind the proof of Theorem 2.6. This implies that on average we change the weights of $O(L / \epsilon)=O\left(\log n / \epsilon^{2}\right)$ edges per update in $G$, for some sufficiently large constants $\alpha, \beta$.
Theorem 2.6. Fix two sufficiently large constants $\alpha, \beta$. In the beginning, when $G$ is an empty graph, initialize a counter COUNT $\leftarrow 0$. Subsequently, each time we change the weight of an already existing edge in the hierarchical partition, set Count $\leftarrow$ Count +1 . Then Count $=O(t L / \epsilon)$ just after we handle the $t^{\text {th }}$ update in $G$.

Define the level of an edge $(y, z)$ to be $\ell(y, z)=\max (\ell(y), \ell(z))$, and note that the weight $w(y, z)$ decreases (resp. increases) iff the edge's level $\ell(y, z)$ goes up (resp. down). There is a potential associated with both nodes and edges. Note that we use the terms "tokens" and "potential" interchangeably.

Each edge $e$ has exactly $2(L-\ell(e))$ tokens. These tokens are assigned as follows. Whenever a new edge is inserted, it receives $2(L-\ell(e))$ tokens. When $e$ moves up a level, it gives one token to each endpoint. Whenever $e$ is deleted, it gives one token to each endpoint. Whenever $e$ moves down a level because one endpoint, say $v$, moves down a level, $e$ receives two tokens from $v$.

Initially and whenever a node moves a level higher, it has no tokens. Whenever a node $v$ moves up a level, only its adjacent edges to the same or lower levels have to be updated as their level changes. Recall that each such edge gives 1 token to $v$, which in turn uses this token to pay for updating the edge. Whenever a node $v$ moves down a level, say from $k$ to $k-1$, it has at most $\beta^{k}$ adjacent edges at level $k$ or below. These are all the edges whose level needs to be updated (which costs a token) and whose potential needs to be increased by two tokens. In this case, we show that $v$ has a number $W_{v}$ of tokens and $W_{v}$ is large enough (i) to pay for the work involved in the update, (ii) to give two tokens to each of the at most $\beta^{k}$ adjacent edges, and (iii) to still have a sufficient number of tokens for being on level $k-1$.

Whenever the level of a node $v$ is not modified but its weight $W_{v}$ decreases because the weight of the adjacent edge $(u, v)$ decreases, the level of $(u, v)$ must have increased and $(u, v)$ gives one token to $v$ (the other one goes to $u$ ). Note that this implies that a change in $W_{v}$ by at most $\beta^{-\ell(v)}$ increases the potential of $v$ by 1 , i.e., the "conversion rate" between weight changes and token changes is $\beta^{\ell(v)}$. Whenever the level of $v$ does not change but its weight $W_{v}$ increases as the level of $(u, v)$ has decreased, no tokens are transferred between $v$ and $(u, v)$. (Technically the potential of $v$ might fall slightly but the change might be so small that we ignore it.) Formally we achieve these potential function changes by setting the potential of every node in $V_{0}$ to 0 and for every other node to $\beta^{\ell(v)} \cdot \max \left(0, \alpha-W_{v}\right)$.

Thus, the crucial claim is that a node $v$ that moves down to level $k-1$ has accumulated a sufficient number $W_{v}$ of tokens, i.e., $W_{v} \geq 3 \beta^{k}+\beta^{k-1} \max \left(0, \alpha-W_{v}(k-1)\right)=X$ (say). We prove this claim by considering two possible cases.
Case 1: Assume first that $v$ was at level $k-1$ immediately before being at level $k$. Recall that, by the definition of the potential function, $v$ had no tokens when it moved up from level $k-1$. However, in this case we know that $W_{v}$ was least $\alpha \beta$ when $v$ moved up and, thus, after adjusting the weights of its adjacent edges to the level change, $W_{v}$ was still at least $\alpha$ after the level change to level $k$. Node $v$ only drops to level $k-1$ if $W_{v}<1$, i.e., while being on level $k$ its weight must have dropped by at least $\alpha-1$. By the above "conversion rate" between weight and tokens this means that $v$ must have received at least $\beta^{k}(\alpha-1)$ tokens from its adjacent edges while it was on level $k$, which is at least $X$ for large enough $\alpha$.
Case 2: Assume next that node $v$ was at level $k+1$ immediately before level $k$. Right after dropping from level $k+1$ node $v$ owned $\beta^{k}\left(\alpha-W_{v}(k)\right)$ tokens. As $v$ has not changed levels since, it did not have to

[^1]give any tokens to edges and did not have to pay for any updates of its adjacent edges. Instead it might have received some tokens from inserted or deleted adjacent edges. Thus, it still owns at least $\beta^{k}\left(\alpha-W_{v}(k)\right)$ tokens. As $W_{v}(k) \leq W_{v}(k-1)$ and $W_{v}(k)<1$ when $v$ drops to level $k-1$, this number of tokens is at least $X$ for $\beta \geq 2$ and $\alpha \geq 3 \beta+1$.

To summarize, whenever an edge is inserted it receives a sufficient number of tokens to pay the cost of future upwards level changes, but also to give a token to its endpoints every time its level increases. These tokens accumulated at the endpoints are sufficient to pay for level decreases of these endpoints because (a) nodes move up to a level when their weight on the new level is at least $\alpha>1$ but only move down when their weight falls below 1 and (b) the weight of edges on the same and lower levels drops by a factor of $\beta$ between two adjacent levels. Thus $\beta^{k}(\alpha-1)$ many edge deletions or edge weight decreases of edges adjacent to node $v$ are necessary to cause $v$ to drop from level $k$ to level $k-1$ (each giving one token to $v$ ), while there are only $\beta^{k-1}$ many edges on levels below $k$ that need to be updated when $v$ drops. Thus, the cost of $v$ 's level drop is $\beta^{k-1}$ and the new potential needed for $v$ on level $k-1$ is $\beta^{k-1}(\alpha-1)$, but $v$ has collected at least $\beta^{k}(\alpha-1)$ tokens, which, by suitable choice of $\beta$ and $\alpha$, is sufficient.

### 2.5 Implementation details

The pseudo-codes for our algorithm, which maintains an $(\alpha, \beta)$-partition that satisfies Invariant 2.4, are given in Figures 2-11.

```
01. }i\leftarrow\operatorname{LEVEL}[\mp@subsup{u}{}{\prime}],j\leftarrow\operatorname{LEVEL}[\mp@subsup{v}{}{\prime}]
02. Call the subroutine UPDATE-LISTS-INSERT ( }\mp@subsup{u}{}{\prime},\mp@subsup{v}{}{\prime})\mathrm{ . See Figure 8.
03. Weight [u']}\leftarrow\operatorname{Weight [ ' }\mp@subsup{u}{}{\prime}]+\mp@subsup{\beta}{}{-max(i,j)}
04. Weight [ [ ' ] \leftarrow Weight [ v}]+\mp@subsup{\beta}{}{-max}(i,j)
05. Call the subroutine UPDATE-STATUS(u'). See Figure 7.
06. Call the subroutine UPDATE-STATUS}(\mp@subsup{v}{}{\prime}). See Figure 7.
07. WHILE the list DIRTY-NODES is nonempty:
08. Let v be the first node in the list DIRTY-NODES.
09. Call the subroutine FIX(v). See Figure 4.
```

Figure 2: INSERT-EDGE $\left(u^{\prime}, v^{\prime}\right)$. It updates the data structures upon insertion of the edge $(u, v)$.

```
01. }\quadi\leftarrow\operatorname{LEVEL}[\mp@subsup{u}{}{\prime}],\quadj\leftarrow\operatorname{LEVEL}[\mp@subsup{v}{}{\prime}]
02. Call the subroutine UPDATE-LISTS-DELETE ( }\mp@subsup{u}{}{\prime},\mp@subsup{v}{}{\prime})\mathrm{ . See Figure 9.
03. WEIGHT [u']}\leftarrow\textrm{WEIGHT}[\mp@subsup{u}{}{\prime}]-\mp@subsup{\beta}{}{-\operatorname{max}(i,j)}\mathrm{ .
04. Weight [v'}]\leftarrow\textrm{WEIGHT}[\mp@subsup{v}{}{\prime}]-\mp@subsup{\beta}{}{-\operatorname{max}(i,j)}\mathrm{ .
05. Call the subroutine UPDATE-STATUS}(\mp@subsup{u}{}{\prime}). See Figure 7.
06. Call the subroutine UPDATE-STATUS( (v'). See Figure 7.
07. While the list Dirty-nodes is nonempty:
08. Let v}\mathrm{ be the first node in the list Dirty-nodes.
09. Call the subroutine FIX(v). See Figure 4.
```

Figure 3: DELETE-EDGE $\left(u^{\prime}, v^{\prime}\right)$. It updates the data structures upon deletion of the edge $(u, v)$.

Initializing the data structures. In the beginning, the graph $G$ has zero edges, all of its nodes are at level zero, and no node is dirty. At that moment, we ensure that our data structures reflect these conditions.

Handling the insertion of an edge. When an edge $\left(u^{\prime}, v^{\prime}\right)$ is inserted into the graph, we implement the procedure in Figure 2. In Step 02, we call the subroutine UPDATE-LISTS-INSERT ( $u^{\prime}, v^{\prime}$ ) to update the neighborhood lists of the nodes $u^{\prime}$ and $v^{\prime}$. In Steps 03-04, we update the weights of $u^{\prime}$ and $v^{\prime}$. In Step 05, we call the subroutine UPDATE-STATUS $\left(u^{\prime}\right)$, which checks if the new weight of $u^{\prime}$ satisfies Invariant 2.4, and accordingly, it updates the bit Status $\left[u^{\prime}\right]$ and the occurrence of $u^{\prime}$ in the list Dirty-nodes. In Step 06 , we perform exactly the same operations on the node $v^{\prime}$.

Every node satisfies Invariant 2.4 at the beginning of the procedure. However, at the end of Step 06, one or both of the nodes in $\left\{u^{\prime}, v^{\prime}\right\}$ can become dirty. Thus, we run the While loop in Steps 07-09 till every node in $G$ becomes clean again. In one iteration of the While loop, we pick any dirty node $v$ and call the subroutine $\operatorname{FIX}(v)$, which either increments or decrements the level of $v$ depending on its weight (Figure 4). The execution of FIX $(v)$ can result in more nodes becoming dirty. So the While loop can potentially have a large number of iterations. To simply the analysis, we charge the runtime of each iteration to the corresponding call to FIX(.). We separately bound the total time taken by all the calls to FIX(.) in Section 2.6.1.

Lemma 2.7. Ignoring the time taken by the calls to FIX(.), an edge-insertion can be handled in $\Theta(1)$ time.
Handling the deletion of an edge. When an edge $\left(u^{\prime}, v^{\prime}\right)$ is deleted from the graph, we implement the procedure in Figure 3, which is very similar to one described above. It first makes all the weights and the neighborhood lists reflect the deletion of the edge. This might lead to one or both the endpoints $\left\{u^{\prime}, v^{\prime}\right\}$ becoming dirty, in which case the procedure keeps on calling the subroutine FIX(.) till every node becomes clean again. As before, we charge the runtime of the WhiLE loop in Steps 07-09 (Figure 3) to the respective calls to FIX(.). We analyze the total runtime of all the calls to FIX(.) in Section 2.6.1.

Lemma 2.8. Ignoring the time taken by the calls to $F I X($.$) , an edge-deletion can be handled in \Theta(1)$ time.

```
01. If Weight \([v]>\alpha \beta\), Then // If true, then Level \([v]<L\) by Lemma 2.3.
02. Call the subroutine MOVE-UP \((v)\). See Figure 5.
03. Else if Weight \([v]<1\) and \(\operatorname{Level}[v]>0\), Then
04. Call the subroutine MOVE-DOWN \((v)\). See Figure 6.
```

Figure 4: FIX $(v)$.

The subroutine FIX $(v)$. The procedure is described in Figure 4. It is called only if the node $v$ is dirty. This can happen only if either (a) $W_{v}>\alpha \beta$, or (b) $W_{v}<1$ and $\ell(v)>0$. In the former case, $W_{v}$ needs to be reduced to bring it down to the range $[1, \alpha \beta]$. So we call the subroutine $\operatorname{MOVE-UP}(v)$ which increments the level of $v$ by one unit. In the latter case, $W_{v}$ needs to be raised to bring it up to the range $[1, \alpha \beta]$. So we call the subroutine MOVE-DOWN $(v)$ which decrements the level of $v$ by one unit.
The subroutine MOVE-UP( $v$ ). See Figure 5. Let $k$ be the level of the node $v$ before the call to MOVE$\mathrm{UP}(v)$. The procedure moves the node $v$ up to level $k+1$. This transition only affects the weights of those edges incident to $v$ whose other endpoints are in $\mathcal{N}_{v}(0, k)$. In addition, $\mathcal{N}_{v}(0, k)$ is precisely the set of nodes that should update the occurrences of $v$ in their neighborhood lists. Further, the node $v$ should move all these nodes from the list Neighbors ${ }_{v}[0, k]$ to the list Neighbors ${ }_{v}[0, k+1]$. Finally, we need to check if any of the nodes in $\mathcal{N}_{v}(0, k)$ changes its status from dirty to clean (or the other way round) due to this transition, and accordingly we need to update the list DIRTY-NODES. All these operations are performed by the While loop in Steps $03-08$ (Figure 5) that runs for $\Theta\left(1+D_{v}(0, k)\right)$ time. ${ }^{2}$

[^2]| 01. | $k \leftarrow \operatorname{LEVEL}[v]$. |
| :---: | :---: |
| 02. | $\operatorname{LEVEL}[v] \leftarrow k+1$. |
| 03. | While the list Neighbors $_{v}[0, k]$ is nonempty: |
| 04. | Let $u$ be a node that appears in Neighbors ${ }_{v}[0, k]$. |
| 05. | Call the subroutine UPDATE-LISTS-UP $(u, v, k)$. See Figure 10. |
| 06. | Weight $[v] \leftarrow$ Weight $[v]-\beta^{-k}+\beta^{-(k+1)}$. |
| 07. | $\mathbf{W e i g h t ~}[u] \leftarrow \operatorname{Weight~}[u]-\beta^{-k}+\beta^{-(k+1)}$. |
| 08. | Call the subroutine UPDATE-STATUS $(u)$. See Figure 7. |
| 09. | In constant time, add all the nodes $u \in$ Neighbors $_{v}[k+1]$ to the list Neighbors ${ }_{v}[0, k+1]$ by adjusting the relevant pointers, and convert Neighbors ${ }_{v}[k+1]$ into an empty list. |
| 10. | Call the subroutine UPDATE-STATUS( $v$ ). See Figure 7. |

Figure 5: MOVE-UP $(v)$. It increments the level of the node $v$ by one.

The nodes in $\mathcal{N}_{v}(k+1)$ should be moved from the list $\operatorname{NeIGHBORS}_{v}[k+1]$ to the list Neighbors ${ }_{v}[0, k+$ 1]. Unlike the nodes in $\mathcal{N}_{v}(0, k)$, however, the nodes in $\mathcal{N}_{v}(k+1)$ need not themselves update the occurrences of $v$ in their neighborhood lists. Due to this reason, we can make the necessary changes with regard to these nodes in constant time, as described in Step 09 (see Figure 5). Finally, in Step 10 we check if the node $v$ satisfies Invariant 2.4 after these transformations. If yes, then we remove the node from the list Dirty-nodes. If no, then the node's level needs to be increased further. It remains dirty and is handled in the next call to $\operatorname{FIX}(v)$. The runtime of $\operatorname{MOVE-UP}(v)$ is summarized in the following lemma.

Lemma 2.9. Let $k$ be the level of the node $v$ just prior to a call to $\operatorname{MOVE-UP(v).~Then~MOVE-UP(v)~runs~}$ for $\Theta\left(1+D_{v}(0, k)\right)$ time.

```
01. }k\leftarrow\operatorname{LEVEL}[v]
02. LEVEL[v]}\leftarrow\operatorname{LEVEL}[v]-1
03. While the list Neighbors
04. Let u be a node that appears in Neighbors }\mp@subsup{v}{v}{}[0,k]\mathrm{ .
05. Call the subroutine UPDATE-LISTS-DOWN(u,v,k). See Figure 11.
06. IF LEVEL}[u]<k,\mathrm{ THEN
07. Weight [u]\leftarrow Weight [u]-\mp@subsup{\beta}{}{-k}+\mp@subsup{\beta}{}{-(k-1)}.
08. Weight [v]}\leftarrow\textrm{WEIGHT}[v]-\mp@subsup{\beta}{}{-k}+\mp@subsup{\beta}{}{-(k-1)}
09.
                                Call the subroutine UPDATE-STATUS(u). See Figure 7.
    Call the subroutine UPDATE-STATUS(v). See Figure 7.
```

Figure 6: MOVE-DOWN $(v)$. It decrements the level of the node $v$ by one.

The subroutine MOVE-DOWN $(v)$. See Figure 6 . Let $k$ be the level of $v$ before the call to MOVE$\operatorname{DOWN}(v)$. The procedure moves the node $v$ down to level $k-1$. This transition only affects the weights of those edges incident to $v$ whose other endpoints are in $\mathcal{N}_{v}(0, k-1)$. In addition, $\mathcal{N}_{v}(0, k-1)$ is precisely the set of nodes that should update the occurrences of $v$ in their neighborhood lists and that might have to change their statuses from dirty to clean (or the other way round). Finally, the node $v$ should move every node in $\mathcal{N}_{v}(0, k)$ from the list $\operatorname{Neighbors~}_{v}[0, k]$ to the list Neighbors ${ }_{v}[0, k-1]$. All these operations are performed by the WHILE loop in Steps 03-09 (see Figure 5) that runs for $\Theta\left(1+D_{v}(0, k)\right)$ time.

In Step 10 (see Figure 6) we check if the node $v$ satisfies Invariant 2.4 under the new circumstances. If yes, then we remove the node from the list Dirty-nodes. If no, then the node's level needs to be decreased further. It remains dirty and is handled in the next call to $\operatorname{FIX}(v)$.

Lemma 2.10. Let $k$ be the level of the node $v$ just prior to a call to $\operatorname{MOVE-DOWN(v)\text {.Thenthecalltothe}}$ subroutine MOVE-DOWN(v) takes $\Theta\left(1+D_{v}(0, k)\right)$ time.

| 01. | IF $1 \leq$ Weight $[u] \leq \alpha \beta$, Then |
| :---: | :---: |
| 02. | If Status $[u]=$ dirty, Then |
| 03. | Remove the node $u$ from the list Dirty-nodes. |
| 04. | Status $[u] \leftarrow$ clean. |
| 05. | Else if Weight $[u]>\alpha \beta$, Then |
| 06. | If Status $[u]=$ clean, Then |
| 07. | Insert the node into the list Dirty-nodes. |
| 08. | Status $[u] \leftarrow$ dirty . |
| 09. | ElSe if Weight $[u]<1$ and Level $[u]>0$, Then |
| 10. | If Status $[u]=$ clean, Then |
| 11. | Insert the node into the list Dirty-nodes. |
| 12. | Status $[u] \leftarrow$ dirty . |
| 13. | ElSe if Weight $[u]<1$ and Level $[u]=0$, Then |
| 14. | If Status $[u]=$ dirty, Then |
| 15. | Remove the node $u$ from the list Dirty-nodes. |
| 16. | Status $[u] \leftarrow$ clean. |

Figure 7: UPDATE-STATUS $(u)$. The subroutine verifies if the node $u$ is dirty or not, and accordingly, it updates the bit Status $[u]$ and the occurrence of the node $u$ in the list Dirty-nodes.

The subroutine UPDATE-STATUS $(v)$. The procedure is described in Figure 7. At every stage of our algorithm, it ensures that the list DIRTY-nODES contains exactly those nodes that violate Invariant 2.4.

| 01. | $i \leftarrow \operatorname{Level}\left[u^{\prime}\right], j \leftarrow \operatorname{LEVEL}\left[v^{\prime}\right]$. |
| :---: | :---: |
| 02. | IF ( $i=j$ ), THEN |
| 03. | Insert node $u^{\prime}$ into the list Neighbors ${ }_{v^{\prime}}[0, i]$. |
| 04. | Insert node $v^{\prime}$ into the list Neighbors ${ }_{u^{\prime}}[0, i]$. |
| 05. | Else if ( $i>j$ ), Then |
| 06. | Insert node $u^{\prime}$ into the list Neighbors ${ }_{v^{\prime}}[i]$. |
| 07. | Insert node $v^{\prime}$ into the list Neighbors ${ }_{u^{\prime}}[0, i]$. |
| 08. | Else if ( $i<j$ ), Then |
| 09. | Insert node $u^{\prime}$ into the list Neighbors ${ }_{v^{\prime}}[0, j]$. |
| 10. | Insert node $v^{\prime}$ into the list Neighbors ${ }_{u^{\prime}}[j]$. |

Figure 8: UPDATE-LISTS-INSERT $\left(u^{\prime}, v^{\prime}\right)$. The edge $\left(u^{\prime}, v^{\prime}\right)$ has been inserted into the graph $G=(V, E)$. The subroutine ensures that the neighborhood lists of $u^{\prime}, v^{\prime}$ reflect this change.

Answering a query. We can handle three types of queries.

1. Is the node $v \in V$ part of the vertex cover?

If Weight $[v] \geq 1$, then we output Yes, else we output No. The query time is $\Theta(1)$.

| 01 | $i \leftarrow \operatorname{LEVEL}\left[u^{\prime}\right], \quad j \leftarrow \operatorname{LEVEL}\left[v^{\prime}\right]$. |
| :---: | :---: |
| 02. | IF ( $i=j$ ), THEN |
| 03. | Delete node $u^{\prime}$ from the list Neighbors ${ }_{v^{\prime}}[0, i]$. |
| 04. | Delete node $v^{\prime}$ from the list Neighbors ${ }_{u^{\prime}}[0, i]$. |
| 05. | Else if ( $i>j$ ), Then |
| 06. | Delete node $u^{\prime}$ from the list Neighbors $_{v^{\prime}}[i]$. |
| 07. | Delete node $v^{\prime}$ from the list Neighbors ${ }_{u^{\prime}}[0, i]$. |
| 08. | Else if ( $i<j$ ), Then |
| 09. | Delete node $u^{\prime}$ from the list Neighbors ${ }_{v^{\prime}}[0, j]$. |
| 10. | Delete node $v^{\prime}$ from the list $\mathrm{Neighbors}_{u^{\prime}}[j]$. |

Figure 9: UPDATE-LISTS-DELETE $\left(u^{\prime}, v^{\prime}\right)$. The edge $\left(u^{\prime}, v^{\prime}\right)$ has been deleted from the graph $G=(V, E)$. The subroutine ensures that the neighborhood lists of $u^{\prime}, v^{\prime}$ reflect this change.

> 01. Delete the node $u$ from the list Neighbors ${ }_{v}[0, k]$.
> 02. Insert the node $u$ into the list Neighbors ${ }_{v}[0, k+1]$.
> 03. IF Level $[u]=k$, Then
> 04. Delete the node $v$ from the list Neighbors ${ }_{u}[0, k]$.
> 05. ElSE IF LEVEl $[u]<k$, Then
> 06. Delete the node $v$ from the list Neighbors $_{u}[k]$.
> 07. Insert the node $v$ into the list Neighbors $_{u}[k+1]$.

Figure 10: UPDATE-LISTS-UP $(u, v, k)$. The level of the node $v$ has changed from $k$ to $k+1$, and we have $u \in \mathcal{N}_{v}(0, k)$. The subroutine ensures that the neighborhood lists of $u, v$ reflect this change.

| 01. | Delete the node $u$ from the |
| :---: | :---: |
| 02. | $\operatorname{IF} \operatorname{Level}[u]=k$, Then |
| 03. | Insert the node $u$ into the list $\mathrm{Neighbors}_{v}[k]$. |
| 04. | Else if Level $[u]=k-1$, Then |
| 05. | Insert the node $u$ into the list Neighbors ${ }_{v}[0, k-1]$. |
| 06. | Delete the node $v$ from the list $\mathrm{Neighbors}_{u}[k]$. |
| 07. | Insert the node $v$ into the list $\mathrm{Neighbors}_{u}[0, k-1]$. |
| 08. | Else if Level $[u]<k-1$, Then |
| 09. | Insert the node $u$ into the list $\mathrm{Neighbors}_{v}[0, k-1]$. |
| 10. | Delete the node $v$ from the list $\mathrm{Neighbors}_{u}[k]$. |
| 11. | Insert the node $v$ into the list $\mathrm{Neighbors}_{u}[k-1]$. |

Figure 11: UPDATE-LISTS-DOWN $(u, v, k)$. The level of the node $v$ has changed from $k$ to $k-1$, and we have $u \in \mathcal{N}_{v}(0, k)$. The subroutine ensures that the neighborhood lists of $u, v$ reflect this change.
2. What is the size of the current vertex cover?

Without increasing the asymptotic runtime of our algorithm, we can easily modify it and maintain an additional counter that will store the number of nodes $v \in V$ with WEIGHT $[v] \geq 1$. To answer this query, we simply return the value of this counter. The query time is $\Theta(1)$.
3. Return all the nodes that are part of the current vertex cover.

Without increasing the asymptotic runtime of our algorithm, we can easily modify it and maintain an additional doubly linked list that will store all the nodes $v \in V$ with WEIGHT $[v] \geq 1$. To answer this query, we simply go through this list and return the nodes that appear in it. The query time is $O\left(\left|V^{*}\right|\right)$, where $V^{*}$ is the vertex cover maintained by our algorithm.

### 2.6 Bounding the amortized update time: Detailed analysis

We devote Section 2.6 to the proof of the following theorem, which bounds the update time of our algorithm. Throughout this section, we set $\alpha \leftarrow 1+3 \epsilon$ and $\beta \leftarrow 1+\epsilon$, where $\epsilon \in(0,1]$ is a small positive constant.

Theorem 2.11. The algorithm in Section 2.5 maintains an $(\alpha, \beta)$-partition that satisfies Invariant 2.4. For every $\epsilon \in(0,1], \alpha=1+3 \epsilon$ and $\beta=1+\epsilon$, the algorithm takes $O\left((t / \epsilon) \log _{1+\epsilon} n\right)$ time to handle $t$ edge updates starting from an empty graph.

Consider the following thought experiment. We have a bank account, and initially, when there are no edges in the graph, the bank account has zero balance. For each subsequent edge insertion/deletion, at most $20 L / \epsilon$ dollars are deposited to the bank account; and for each unit of computation performed by our algorithm, at least one dollar is withdrawn from it. We show that the bank account never runs out of money, and this gives a running time bound of $O(t L / \epsilon)=O\left((t / \epsilon) \log _{\beta}(n / \alpha)\right)=O\left((t / \epsilon) \log _{1+\epsilon} n\right)$ for handling $t$ edge updates starting from an empty graph.

Let $\mathcal{B}$ denote the total amount of money (or potential) in the bank account at the present moment. We keep track of $\mathcal{B}$ by distributing an $\epsilon$-fraction of it among the nodes and the current set of edges in the graph.

$$
\begin{equation*}
\mathcal{B}=(1 / \epsilon) \cdot\left(\sum_{e \in E} \Phi(e)+\sum_{v \in V} \Psi(v)\right) \tag{12}
\end{equation*}
$$

In the above equation, the amount of money (or potential) associated with an edge $e \in E$ is given by $\Phi(e)$, and the amount of money (or potential) associated with a node $v \in V$ is given by $\Psi(v)$. To ease notation, for each edge $e=(u, v) \in E$, we use the symbols $\Phi(e), \Phi(u, v)$ and $\Phi(v, u)$ interchangeably.

We call a node $v \in V$ passive if we have $\mathcal{N}_{v}=\emptyset$ throughout the duration of a time interval that starts at the beginning of the algorithm (when $E=\emptyset$ ) and ends at the present moment. Let $V_{\text {passive }} \subseteq V$ denote the set of all nodes that are currently passive.

At every point in time, all the potentials $\Phi(u, v), \Psi(v)$ are determined by the two invariants stated below.
Invariant 2.12. For every (unordered) pair of node $\{u, v\}, u, v \in V$, we have:

$$
\Phi(u, v)= \begin{cases}(1+\epsilon) \cdot(L-\max (\ell(u), \ell(v))) & \text { if }(u, v) \in E \\ 0 & \text { if }(u, v) \notin E\end{cases}
$$

Invariant 2.13. For every node $v \in V$, we have:

$$
\Psi(v)= \begin{cases}\epsilon \cdot(L-\ell(v))+\left(\beta^{\ell(v)+1} /(\beta-1)\right) \cdot \max \left(0, \alpha-W_{v}\right) & \text { if } v \notin V_{\text {passive }} \\ 0 & \text { if } v \in V_{\text {passive }}\end{cases}
$$

Initialization. When the algorithm starts, the graph has zero edges, all the nodes are at level 0 , and every node is passive. At that moment, Invariant 2.13 sets $\Psi(v)=0$ for all nodes $v \in V$. Consequently, equation 12 implies that the potential $\mathcal{B}$ is also set to zero. This is consistent with our requirement that initially the bank account ought to have zero balance.

Insertion of an edge. By Lemma 2.7, the time taken to handle an edge insertion, ignoring the calls to the FIX(.) subroutine, is $\Theta(1)$. According to our framework, we are allowed to deposit at most $20 L / \epsilon$ dollars to the bank account, and a withdrawal of one dollar from the same place is sufficient to pay for the computation performed. Thus, due to the Steps 01-06 in Figure 2, the net increase in the potential $\mathcal{B}$ ought to be no more than $20 L / \epsilon-1$. We show below that this is indeed the case.

When an edge $(u, v)$ is inserted into the graph, $\Phi(u, v)$ increases by at most $(1+\epsilon) L$. Next, we consider two possible scenarios to bound the increase in $\Psi(v)$.

1. The node $v$ was passive prior to the insertion of the edge. Clearly, in this case the node is at level zero, and $\Psi(v)=0$ before the insertion. After the insertion the node is not passive anymore, and we have $\Psi(v) \leq \epsilon L+\alpha \beta /(\beta-1)=\epsilon L+(1+3 \epsilon)(1+\epsilon) / \epsilon \leq 9 L$. Since $\epsilon \in[0,1]$ and $L=\log _{\beta}(n / \alpha)=$ $\log _{1+\epsilon}(n /(1+3 \epsilon))$, the last inequality holds as long as $n \geq 20$. To summarize, we conclude that $\Psi(v)$ increases by at most $9 L$.
2. The node $v$ was not passive prior to the insertion of the edge. Clearly, in this case the node $v$ is also not passive afterwards. The weight $W_{v}$ increases and the quantity $\max \left(0, \alpha-W_{v}\right)$ decreases due to the insertion, which means that the potential $\Psi(v)$ decreases.

Similarly, we conclude that either $\Psi(u)$ increases by at most $9 L$ or it actually decreases. The potentials of the remaining nodes and edges do not change. Hence, by equation 12 , the net increase in $\mathcal{B}$ is at most $18 L / \epsilon$.
Deletion of an edge. The analysis is very similar to the one described above. By Lemma 2.8, the time taken to handle an edge deletion, ignoring the calls to FIX(.), is $\Theta(1)$. According to our framework, we are allowed to deposit at most $20 L / \epsilon$ dollars to the bank account, and a withdrawal of one dollar from the same place is sufficient to pay for the computation performed. Thus, due to the Steps 01-06 in Figure 3, the net increase in the potential $\mathcal{B}$ ought to be no more than $20 L / \epsilon-1$. We show below that this is indeed the case.

When an edge $(u, v)$ is deleted from the graph, the potential $\Phi(u, v)$ decreases. Next, note that the weight $W_{v}$ decreases by at most $\beta^{-\ell(v)}$, and so the quantity $\max \left(0, \alpha-W_{v}\right)$ increases by at most $\beta^{-\ell(v)}$. As the node $v$ was not passive before the edge deletion, it is also not passive afterwards. Thus, $\Psi(v)$ increases by at most $\left(\beta^{\ell(v)+1} /(\beta-1)\right) \cdot \beta^{-\ell(v)}=2 / \epsilon \leq 2 L$. The last equality holds as long as $n \geq 20$. We similarly conclude that $\Psi(u)$ increases by at most $2 L$. The potentials of the remaining nodes and edges do not change. Hence, by equation 12 , the net increase in $\mathcal{B}$ is at most $4 L / \epsilon$.

It remains to analyze the total runtime of all the calls to FIX(.). This is done in Section 2.6.1.

### 2.6.1 Analysis of the FIX(v) subroutine

We analyze a single call to the subroutine FIX $(v)$. Throughout this section, we use the superscript 0 (resp. 1) on a symbol to denote its state at the time instant immediately prior to (resp. after) the execution of FIX $(v)$. Further, we preface a symbol with $\delta$ to denote the net decrease in its value due to the call to FIX $(v)$. For example, consider the potential $\mathcal{B}$. We have $\mathcal{B}=\mathcal{B}^{0}$ immediately before $\operatorname{FIX}(v)$ is called, and $\mathcal{B}=\mathcal{B}^{1}$ immediately after $\operatorname{FIX}(v)$ terminates. We also have $\delta \mathcal{B}=\mathcal{B}^{0}-\mathcal{B}^{1}$. We will prove the following theorem.

Theorem 2.14. We have $\delta \mathcal{B}>0$ and $\delta \mathcal{B} \geq T$, where $T$ denotes the runtime of $F I X(v)$. In other words, the money withdrawn from the bank account during the execution of $\operatorname{FIX}(v)$ is sufficient to pay for the computation performed by $\operatorname{FIX}(v)$.

The call to FIX $(v)$ affects only the potentials of the nodes $u \in \mathcal{N}_{v} \cup\{v\}$ and that of the edges $e \in$ $\left\{(u, v): u \in \mathcal{N}_{v}\right\}$. This observation, coupled with equation 12 , gives us the following guarantee.

$$
\begin{equation*}
\delta \mathcal{B}=(1 / \epsilon) \cdot\left(\delta \Psi(v)+\sum_{u \in \mathcal{N}_{v}} \delta \Phi(u, v)+\sum_{u \in \mathcal{N}_{v}} \delta \Psi(u)\right) \tag{13}
\end{equation*}
$$

The call to FIX $(v)$ does not change (a) the neighborhood structure of the node $v$, and (b) the level and the overall degree of any node $u \neq v$.

$$
\begin{align*}
\mathcal{N}_{v}^{0}(i) & =\mathcal{N}_{v}^{1}(i) \text { for all } i \in\{0, \ldots, L\} .  \tag{14}\\
\ell^{0}(u) & =\ell^{1}(u) \text { for all } u \in V \backslash\{v\} .  \tag{15}\\
D_{u}^{0} & =D_{u}^{1} \text { for all } u \in V \backslash\{v\} . \tag{16}
\end{align*}
$$

Accordingly, to ease notation we do not put any superscript on the following symbols, as the quantities they refer to remain the same throughout the duration of FIX $(v)$.

$$
\begin{cases}\mathcal{N}_{v}, D_{v} & \\ \mathcal{N}_{v}(i), D_{v}(i), W_{v}(i) & \text { for all } i \in\{0, \ldots, L\} \\ \mathcal{N}_{v}(i, j), D_{v}(i, j) & \text { for all } i, j \in\{0, \ldots, L\}, i \leq j \\ \ell(u), D_{u} & \text { for all } u \in V \backslash\{v\} .\end{cases}
$$

Since the node $v$ is dirty when the subroutine FIX $(v)$ is called, it follows that neither the node $v$ nor the nodes $u \in \mathcal{N}_{v}$ are passive. ${ }^{3}$ Applying Invariant 2.13, we get:

$$
\begin{equation*}
\Psi(u)=\epsilon \cdot(L-\ell(u))+\left(\beta^{\ell(u)+1} /(\beta-1)\right) \cdot \max \left(0, \alpha-W_{u}\right) \quad \text { for all nodes } u \in \mathcal{N}_{v} \cup\{v\} . \tag{17}
\end{equation*}
$$

We divide the proof of Theorem 2.14 into two possible cases, depending upon whether the call to FIX $(v)$ increments or decrements the level of $v$. The main approach to the proof remains the same in each case. We first give an upper bound on the running time $T$. Next, we separately lower bound each of the following quantities: $\delta \Psi(v), \delta \Phi(u, v)$ for all $u \in \mathcal{N}_{v}$, and $\delta \Psi(u)$ for all $u \in \mathcal{N}_{v}$. Finally, applying equation 13, we derive that $\delta \mathcal{B} \geq T$.
Case 1: The subroutine $\operatorname{FIX}(v)$ increases the level of the node $v$ from $k$ to $(k+1)$.
Lemma 2.15. We have $T \leq 1+D_{v}(0, k)$.
Proof. FIX $(v)$ calls the subroutine MOVE-UP $(v)$, which, by Lemma 2.9, runs for $\Theta\left(1+D_{v}(0, k)\right)$ time. ${ }^{4}$

Lemma 2.16. We have $\delta \Psi(v)=\epsilon$.
Proof. The subroutine FIX $(v)$ calls the subroutine $\operatorname{MOVE-UP}(v, j)$ in Step 02 (Figure 4). Hence, Step 01 (Figure 4) guarantees that $W_{v}^{0}=W_{v}(k)>\alpha \beta$. Next, from Lemma 2.3 we infer that $W_{v}^{1}=W_{v}(k+1) \geq$ $\beta^{-1} \cdot W_{v}(k)>\alpha$. Since both $W_{v}^{0}, W_{v}^{1}>\alpha$, we get:

$$
\begin{aligned}
& \Psi^{0}(v)=\epsilon \cdot(L-k)+\left(\beta^{k+1} /(\beta-1)\right) \cdot \max \left(0, \alpha-W_{v}^{0}\right)=\epsilon \cdot(L-k) \\
& \Psi^{1}(v)=\epsilon \cdot(L-k-1)+\left(\beta^{k+2} /(\beta-1)\right) \cdot \max \left(0, \alpha-W_{v}^{1}\right)=\epsilon \cdot(L-k-1)
\end{aligned}
$$

It follows that $\delta \Psi(v)=\Psi^{0}(v)-\Psi^{1}(v)=\epsilon$.

[^3]Lemma 2.17. For every node $u \in \mathcal{N}_{v}$, we have:

$$
\delta \Phi(u, v)= \begin{cases}(1+\epsilon) & \text { if } u \in \mathcal{N}_{v}(0, k) \\ 0 & \text { if } u \in \mathcal{N}_{v}(k+1, L) .\end{cases}
$$

Proof. If $u \in \mathcal{N}_{v}(0, k)$, then we have $\Phi^{0}(u, v)=(1+\epsilon) \cdot(L-k)$ and $\Phi^{1}(u, v)=(1+\epsilon) \cdot(L-k-1)$. It follows that $\delta \Phi(u, v)=\Phi^{0}(u, v)-\Phi^{1}(u, v)=(1+\epsilon)$.

In contrast, if $u \in \mathcal{N}_{v}(k+1, L)$, then we have $\Phi^{0}(u, v)=\Phi^{1}(u, v)=(1+\epsilon) \cdot(L-\ell(u))$. Hence, we get $\delta \Phi(u, v)=\Phi^{0}(u, v)-\Phi^{1}(u, v)=0$.

Lemma 2.18. For every node $u \in \mathcal{N}_{v}$, we have:

$$
\delta \Psi(u) \geq \begin{cases}-1 & \text { if } u \in \mathcal{N}_{v}(0, k) \\ 0 & \text { if } u \in \mathcal{N}_{v}(k+1, L) .\end{cases}
$$

Proof. Consider any node $u \in \mathcal{N}_{v}(k+1, L)$. Since $k<\ell(u)$, we have $w^{0}(u, v)=w^{1}(u, v)$, and this implies that $W_{u}^{0}=W_{u}^{1}$. Thus, we get $\delta \Psi(u)=0$.

Next, fix any node $u \in \mathcal{N}_{v}(0, k)$. Note that $\delta W_{u}=\delta w(u, v)=\beta^{-k}-\beta^{-(k+1)}=(\beta-1) / \beta^{k+1}$. Using this observation, and the fact that $\ell(u) \leq k$, we infer that:

$$
\delta \Psi(u) \geq-\left(\beta^{\ell(u)+1} /(\beta-1)\right) \cdot \delta W_{u}=-\beta^{\ell(u)+1} / \beta^{k+1} \geq-1 .
$$

Proof of Theorem 2.14 (for Case 1). From Lemmas 2.16, 2.17, 2.18 and equation 13, we derive the following bound.

$$
\begin{aligned}
\delta \mathcal{B} & =(1 / \epsilon) \cdot\left(\delta \Psi(v)+\sum_{u \in \mathcal{N}_{v}} \delta \Phi(u, v)+\sum_{u \in \mathcal{N}_{v}} \delta \Psi(u)\right) \\
& \geq(1 / \epsilon) \cdot\left(\epsilon+(1+\epsilon) \cdot D_{v}(0, k)-D_{v}(0, k)\right) \\
& =1+D_{v}(0, k)
\end{aligned}
$$

The theorem (for Case 1) now follows from Lemma 2.15.

## Case 2: The subroutine FIX $(v)$ decreases the level of the node $v$ from $k$ to $k-1$.

Claim 2.19. We have $W_{v}^{0}=W_{v}(k)<1$ and $D_{v}(0, k) \leq \beta^{k}$.
Proof. As FIX $(v)$ calls the subroutine MOVE-DOWN( $v$ ), Step 03 (Figure 4) ensures that $W_{v}^{0}=W_{v}(k)<$ 1 . Since $\ell^{0}(v)=k$, we have $w^{0}(u, v) \geq \beta^{-k}$ for all $u \in \mathcal{N}_{v}$. We conclude that:

$$
1>W_{v}^{0} \geq \sum_{u \in \mathcal{N}_{v}(0, k)} w^{0}(u, v) \geq \beta^{-k} \cdot D_{v}(0, k) .
$$

Thus, we get $D_{v}(0, k) \leq \beta^{k}$.
Lemma 2.20. We have $T \leq \beta^{k}$.
Proof. FIX $(v)$ calls the subroutine MOVE-DOWN $(v)$, which runs for $\Theta\left(1+D_{v}(0, k)\right)$ time by Lemma 2.10. ${ }^{5}$ The lemma now follows from Claim 2.19.

[^4]Lemma 2.21. For every node $u \in \mathcal{N}_{v}$, we have $\delta \Psi(u) \geq 0$.
Proof. Fix any node $u \in \mathcal{N}_{v}$. As the level of the node $v$ decreases from $k$ to $k-1$, we infer that $w^{0}(u, v) \leq$ $w^{1}(u, v)$, and, accordingly, we get $W_{u}^{0} \leq W_{u}^{1}$. Since $\Psi(u)=\epsilon \cdot(L-\ell(u))+\beta^{\ell(u)} \cdot \max \left(0, \alpha-W_{u}\right)$, we derive that $\Psi^{0}(u) \geq \Psi^{1}(u)$. Thus, we have $\delta \Psi(u)=\Psi^{0}(u)-\Psi^{1}(u) \geq 0$.

Lemma 2.22. For every node $u \in \mathcal{N}_{v}$, we have:

$$
\delta \Phi(u, v)= \begin{cases}0 & \text { if } u \in \mathcal{N}_{v}(k, L) ; \\ -(1+\epsilon) & \text { if } u \in \mathcal{N}_{v}(0, k-1) ;\end{cases}
$$

Proof. Fix any node $u \in \mathcal{N}_{v}$. We prove the lemma by considering two possible scenarios.

1. We have $u \in \mathcal{N}_{v}(k, L)$. As FIX $(v)$ decreases the level of the node $v$ from $k$ to $k-1$, we infer that $\Phi^{0}(u, v)=\Phi^{1}(u, v)=(1+\epsilon) \cdot(L-\ell(u))$. Hence, we get $\delta \Phi(u, v)=\Phi^{1}(u, v)-\Phi^{0}(u, v)=0$.
2. We have $u \in \mathcal{N}_{v}(0, k-1)$. Since $\operatorname{FIX}(v)$ decreases the level of node $v$ from $k$ to $k-1$, we infer that $\Phi^{0}(u, v)=(1+\epsilon) \cdot(L-k)$ and $\Phi^{1}(u, v)=(1+\epsilon) \cdot(L-k+1)$. Hence, we get $\delta \Phi(u, v)=$ $\Phi^{1}(u, v)-\Phi^{0}(u, v)=-(1+\epsilon)$.

We partition $W_{v}^{0}$ into two parts: $x$ and $y$. The first part denotes the contributions towards $W_{v}^{0}$ by the neighbors of $v$ that lie below level $k$, while the second part denotes the contribution towards $W_{v}^{0}$ by the neighbors of $v$ that lie on or above level $k$. Thus, we get the following equations.

$$
\begin{align*}
W_{v}^{0} & =x+y \leq 1  \tag{18}\\
x & =\sum_{u \in \mathcal{N}_{v}(0, k-1)} w^{0}(u, v)=\beta^{-k} \cdot D_{v}(0, k-1)  \tag{19}\\
y & =\sum_{u \in \mathcal{N}_{v}(k, L)} w^{0}(u, v) \tag{20}
\end{align*}
$$

Lemma 2.23. We have $\sum_{u \in \mathcal{N}_{v}} \delta \Phi(u, v)=-(1+\epsilon) \cdot x \cdot \beta^{k}$.
Proof. Lemma 2.22 implies that $\sum_{u \in \mathcal{N}_{v}} \delta \Phi(u, v)=-(1+\epsilon) \cdot D_{v}(0, k-1)$. Applying equation 19 , we infer that $D_{v}(0, k-1)=x \cdot \beta^{k}$. The lemma follows.

Lemma 2.24. We have:

$$
\delta \Psi(v)=-\epsilon+(\alpha-x-y) \cdot\left(\beta^{k+1} /(\beta-1)\right)-\max (0, \alpha-\beta x-y) \cdot\left(\beta^{k} /(\beta-1)\right) .
$$

Proof. Equation 18 states that $W_{v}^{0}=x+y<1$. Since $\ell^{0}(v)=k$, we get:

$$
\begin{equation*}
\Psi^{0}(v)=\epsilon(L-k)+(\alpha-x-y) \cdot\left(\beta^{k+1} /(\beta-1)\right) \tag{21}
\end{equation*}
$$

As the node $v$ decreases its level from $k$ to $k-1$, we have:

$$
w^{1}(u, v)= \begin{cases}\beta \cdot w^{0}(u, v) & \text { if } u \in \mathcal{N}_{v}(0, k-1) \\ w^{0}(u, v) & \text { if } u \in \mathcal{N}_{v}(k, L)\end{cases}
$$

Accordingly, we have $W_{v}^{1}=\beta \cdot x+y$, which implies the following equation.

$$
\begin{equation*}
\Psi^{1}(v)=\epsilon(L-k+1)+\max (0, \alpha-\beta x-y) \cdot\left(\beta^{k} /(\beta-1)\right) \tag{22}
\end{equation*}
$$

Since $\delta \Psi(v)=\Psi^{0}(v)-\Psi^{1}(v)$, the lemma now follows from equations 21 and 22 .

Proof of Theorem 2.14 (for Case 2). We consider two possible scenarios depending upon the value of ( $\alpha-$ $\beta x-y$ ). We show that in each case $\delta \mathcal{B} \geq \beta^{k}$. The theorem (for Case 2) then follows from Lemma 2.20.

1. Suppose that $(\alpha-\beta x-y)<0$. From Lemmas 2.21, 2.23, 2.24 and equation 13, we derive:

$$
\begin{align*}
\epsilon \cdot \delta \mathcal{B} & =\sum_{u \in \mathcal{N}_{v}} \delta \Psi(u)+\sum_{u \in \mathcal{N}_{v}} \delta \Phi(u, v)+\delta \Psi(v) \\
& \geq 0-(1+\epsilon) \cdot x \cdot \beta^{k}-\epsilon+(\alpha-x-y) \cdot \beta^{k+1} /(\beta-1) \\
& \geq-\epsilon-(1+\epsilon) \cdot \beta^{k}+(\alpha-1) \cdot \beta^{k+1} /(\beta-1)  \tag{equation18}\\
& =-\epsilon+\frac{\beta^{k}}{(\beta-1)} \cdot\{(\alpha-1) \cdot \beta-(1+\epsilon)(\beta-1)\} \\
& =-\epsilon+2 \cdot(1+\epsilon) \cdot \beta^{k} \\
& \geq \epsilon \cdot \beta^{k}
\end{align*}
$$

$$
=-\epsilon+2 \cdot(1+\epsilon) \cdot \beta^{k} \quad \quad(\text { since } \alpha=1+3 \epsilon \text { and } \beta=1+\epsilon)
$$

(since $\beta>1, \epsilon \leq 1$ )
2. Suppose that $(\alpha-\beta x-y) \geq 0$. From Lemmas 2.21, 2.23, 2.24 and equation 13, we derive:

$$
\begin{aligned}
\epsilon \cdot \delta \mathcal{B} & =\sum_{u \in \mathcal{N}_{v}} \delta \Psi(u)+\sum_{u \in \mathcal{N}_{v}} \delta \Phi(u, v)+\delta \Psi(v) \\
& \geq 0-(1+\epsilon) \cdot x \cdot \beta^{k}-\epsilon+(\alpha-x-y) \cdot \beta^{k+1} /(\beta-1)-(\alpha-\beta x-y) \cdot \beta^{k} /(\beta-1) \\
& =-\epsilon+\frac{\beta^{k}}{(\beta-1)} \cdot\{(\alpha-x-y) \cdot \beta-(1+\epsilon) \cdot x \cdot(\beta-1)-(\alpha-\beta x-y)\} \\
& =-\epsilon+\frac{\beta^{k}}{(\beta-1)} \cdot\{(\alpha-x-y) \cdot(\beta-1)-\epsilon \cdot x \cdot(\beta-1)\} \\
& \geq-\epsilon+\frac{\beta^{k}}{(\beta-1)} \cdot\{(\alpha-1) \cdot(\beta-1)-\epsilon \cdot(\beta-1)\} \quad(\text { since } 0 \leq x+y \leq 1) \\
& =-\epsilon+2 \cdot \epsilon \cdot \beta^{k} \quad(\text { since } \alpha=1+3 \epsilon \text { and } \beta=1+\epsilon) \\
& \geq \epsilon \cdot \beta^{k} \quad(\text { since } \beta>1)
\end{aligned}
$$

## 3 Dynamic Matching: Preliminaries

We are given an input graph $G=(V, E)$ that is being updated dynamically through a sequence of edge insertions/deletions. We want to maintain an approximately maximum matching in $G$. We will present three different algorithms for this problem. They are described in Sections 4, 5 and 6. All these algorithms, however, will use two key ideas.

1. It is easy to maintain a good approximate matching in a bounded degree graph (see Section 3.1).
2. Every graph contains a subgraph of bounded degree, called its kernel, that approximately preserves the size of the maximum matching (see Section 3.2).

In Section 3.3, we give a static algorithm for building a kernel in a graph. In Section 3.4, we present the data structures that will be used in Sections 4, 5 and 6 for maintaining a kernel in a dynamic setting.

Query time. In this paper, all the data structures for dynamic matching explicitly maintain the set of matched edges. Accordingly, using appropriate pointers, we can support the following queries.

- Report the size of the matching $M$ maintained by the data structure in $O(1)$ time.
- Report the edges in $M$ in $O(1)$ time per edge.
- In $O(1)$ time, report whether or not a given edge $(u, v)$ is part of the matching.
- In $O(1)$ time, report whether or not a given node $u$ is matched in $M$, and if yes, in $O(1)$ time report the node it is matched to.


### 3.1 Maintaining an approximate matching in a bounded degree graph.

It is well known that maintaining a maximal matching is easy in a bounded degree graph. Still, for the sake of completeness, we state this formally in the theorem below.

Theorem 3.1. Consider a dynamic graph $\mathcal{G}=(\mathcal{V}, \mathcal{E})$, where the maximum degree of a node is always upper bounded by $\Delta$. Then we have an algorithm for maintaining a maximal matching $M$ in $\mathcal{G}$ that handles each edge insertion in $O(1)$ time and each edge deletion in $O(\Delta)$ time.

Proof. When an edge $(u, v)$ is inserted into the graph $\mathcal{G}$, we check if both the nodes $u, v$ are unmatched in $M$. If yes, then we set $M \leftarrow M \cup\{(u, v)\}$. If no, then we do not change the matching $M$.

When an edge $(u, v) \in M$ is deleted from the graph $\mathcal{G}$, we first set $M \leftarrow M \backslash\{(u, v)\}$. Next, we go through all the neighbors of $u$ to find out if any one of them is unmatched in $M$. If there exists such a node $x$, with $(u, x) \in \mathcal{E} \backslash M$, then we set $M \leftarrow M \cup\{(u, x)\}$. Finally, we perform exactly the same operations on the node $v$.

When an edge $(u, v) \in \mathcal{E} \backslash M$ is deleted from the graph $\mathcal{G}$, we leave the matching $M$ as it is.
Clearly, this algorithm maintains a maximal matching, and handles each edge insertion (resp. deletion) in $O(1)$ (resp. $O(\Delta))$ time.

We now recall the definition of an augmenting path in a graph with respect to a given matching.
Definition 3.2. Consider any graph $\mathcal{G}=(\mathcal{V}, \mathcal{E})$ and any matching $M \subseteq \mathcal{E}$ in $\mathcal{G}$. For any nonnegative integer $k$, an augmenting path of length $2 k+1$ in $\mathcal{G}$, w.r.t. $M$, is a path $\left(v_{1}, \ldots, v_{2 k+1}\right)$ that satisfies the following properties.

- We have $\left(v_{1}, v_{2}\right) \in \mathcal{E} \backslash M$ and $\left(v_{2 k}, v_{2 k+1}\right) \in \mathcal{E} \backslash M$.
- For every $i \in\{1, \ldots, 2 k-1\}$, if $\left(v_{i}, v_{i+1}\right) \in \mathcal{E} \backslash M$, then $\left(v_{i+1}, v_{i+2}\right) \in M$. And if $\left(v_{i}, v_{i+1}\right) \in M$, then $\left(v_{i+1}, v_{i+2}\right) \in \mathcal{E} \backslash M$.

The following theorem is well known (see e.g., [5]).
Theorem 3.3. If $M$ is a matching in a graph $\mathcal{G}$ such that every augmenting path in $\mathcal{G}$ w.r.t. $M$ has length at least $(2 k+1)$, then $M$ is a $(1+1 / k)$-approximate maximum matching in $\mathcal{G}$.

Following the approach from [12], we show in the next theorem shows that in $O(\Delta)$ worst-case update time we can also maintain a $3 / 2$-approximate matching in a graph with maximum degree $\Delta$.

Theorem 3.4. Consider a dynamic graph $\mathcal{G}=(\mathcal{V}, \mathcal{E})$ that is being updated through a sequence of edge insertions/deletions. Furthermore, the maximum degree of a node in $\mathcal{G}$ is always upper bounded by $\Delta$. Then we have a dynamic algorithm for maintaining a matching $M$ in $\mathcal{G}$ that has the following properties.

- The algorithm guarantees that every augmenting path in $\mathcal{G}$ (w.r.t. $M$ ) has length at least five.
- Each edge insertion into $\mathcal{G}$ is handled in $O(\Delta)$ worst-case time.
- Each edge deletion from $\mathcal{G}$ is handled in $O(\Delta)$ worst-case time.


### 3.1. Proof of Theorem 3.4

The algorithm maintains the following data structures for each node $v \in \mathcal{V}$.

- A doubly linked list Neighbors $(v)$. It consists of the set of nodes that are adjacent to $v$ in $\mathcal{G}$.
- A doubly linked list Free-Neighbors $(v)$. It consists of the unmatched (in $M$ ) neighbors of $v$.

Furthermore, the algorithm maintains two $|\mathcal{V}| \times|\mathcal{V}|$ matrices $Q_{1}$ and $Q_{2}$. For every pair of nodes $u, v$, if $u \in \operatorname{Neighbors}(v)$, then the entry $Q_{1}[u, v]$ stores a pointer to the position of $u$ in the list Neighbors $(v)$. Similarly, if $u \in \operatorname{Free-Neighbors}(v)$, then the entry $Q_{2}[u, v]$ stores a pointer to the position of $u$ in the list Free-Neighbors $(v)$. Thus, a node $u$ can be inserted into (or deleted from) a list Neighbors $(v)$ or Free-Neighbors $(v)$ in constant time.
Handling an edge insertion. Suppose that the edge $(u, v)$ is inserted into $\mathcal{G}$. We make our data structures consistent with $\mathcal{G}$ by adding the node $u$ to the list $\operatorname{Neighbors~}(v)$ and the node $v$ to the list Neighbors $(u)$.

Before the insertion, every augmenting path in $\mathcal{G}$ (w.r.t $M$ ) had length at least five. After the insertion, some augmenting paths (containing the edge $(u, v)$ ) may be created that are of length one or three. To resolve such paths, we consider three possible scenarios.

- Case 1. Both the endpoints $u$ and $v$ are unmatched in $M$. Here, we set $M \leftarrow M \cup\{(u, v)\}$.
- Case 2. One the endpoints (say $u$ ) is matched in $M$, while the other endpoint $v$ is unmatched in $M$. Let $x$ be the node $u$ is matched to, i.e., $(u, x) \in M$. If the list $\operatorname{Free} \operatorname{Neighbors}(x)$ is empty, then we do nothing and terminate the procedure. Otherwise, let $y$ be a node that appears in Free-Neighbors $(x)$. Clearly, $(v, u, x, y)$ is an augmenting path of length three. We resolve this path by setting $M \leftarrow(M \backslash\{(u, x)\}) \cup\{(u, v),(x, y)\}$. The nodes $v, y$ are now matched in $M$. To reflect this change, for every node $z \in \operatorname{Neighbors}(y)$, we remove $y$ from the list Free-Neighbors $(z)$. Similarly, for every node $z \in \operatorname{Neighbors}(v)$, we remove $v$ from the list Free-Neighbors $(v)$.
- Case 3. Both the endpoints $u$ and $v$ are matched in $M$. In this case, the insertion of the edge $(u, v)$ does not create any augmenting path in $\mathcal{G}$ (w.r.t $M$ ) of length one or three. Hence, we do nothing and terminate the procedure.

If we are in Case 1 or in Case 3, then the procedure takes $O(1)$ time. In contrast, if we are in Case 2 , then the procedure takes $O(\Delta)$ time. So the worst-case update time for handling an edge insertion is $O(\Delta)$.
Handling an edge deletion. Suppose that the edge $(u, v)$ is deleted from the graph $\mathcal{G}$. We first make our data structures consistent with $\mathcal{G}$ by removing the node $u$ from the list Neighbors $(v)$ and the node $v$ from the list Neighbors $(u)$. Next, we consider two possible scenarios.

- Case 1. $(u, v) \notin M$. In this case, no augmenting path of length one or three has been created. Thus, we do nothing and terminate the procedure.
- Case 2. $(u, v) \in M$. We set $M \leftarrow M \backslash\{(u, v)\}$. The nodes $u$ and $v$ are no longer matched in $M$. To reflect this change, for every node $x \in \operatorname{Neighbors}(u)$, we add $u$ to the list Free-Neighbors $(x)$. Similarly, for every node $x \in \operatorname{Neighbors}(v)$, we add $v$ to the list Free-Neighbors $(x)$. This may create some augmenting paths in $\mathcal{G}$ (w.r.t. $M$ ) that are of length one or three. All such bad augmenting paths, however, start from either the node $u$ or the node $v$. To fix this problem, we call the subroutines Find-mate $(u)$ and $\operatorname{Find}-\operatorname{mate}(v)$, one after the other (see Figure 12).
If the set of unmatched neighbors of $u$ (resp. $v$ ) is nonempty, then FInd-mate $(u)$ (resp. Find-mate $(v)$ ) matches $u$ (resp. $v$ ) to any one of them. If the subroutine $\operatorname{Find}-\operatorname{mate}(u)$ (resp. Find-mate $(v)$ ) succeeds in finding a mate for $u$ (resp. $v$ ), then we know for sure that there is no more augmenting path in $\mathcal{G}$ of length one or three that starts from $u$ (resp. $v$ ). In case of a failure, however, we have to deal with the possibility that there is an augmenting path of length three in the graph $\mathcal{G}$ that starts from $u$ (resp. $v$ ).
To take care of this issue, we finally call the subroutines $\operatorname{Resolve}(u)$ and $\operatorname{Resolve}(v)$, one after the other (see Figure 13). The former (resp. latter) subroutine resolves an augmenting path of length three, if there is any, starting from the node $u$ (resp. $v$ ). At their termination, we are guaranteed that every augmenting path in $\mathcal{G}$ (w.r.t. $M$ ) has length at least five.

For all $x \in \operatorname{NeIGHBORS}(u)$ :
If $x$ is unmatched in $M$, THEN Set $M \leftarrow M \cup\{(u, x)\}$. For all $z \in \operatorname{Neighbors~}(u)$ :

Remove $u$ from the list Free-Neighbors $(z)$. For all $z \in \operatorname{Neighbors}(x)$ :

Remove $x$ from the list Free-Neighbors $(z)$.
Figure 12: Find-mate $(u)$.

```
IF \(u\) is already matched in \(M\), THEN
RETURN.
```

For all $x \in \operatorname{NeIGHBORS}(u)$ :
We must have that $x$ is matched in $M$.
Let $y$ be the node $x$ is matched to, i.e., $(x, y) \in M$.
If Free-Neighbors $(y) \backslash\{u\} \neq \emptyset$, Then
Consider any node $z \in \operatorname{Free-Neighbors}(y) \backslash\{u\}$.
Resolve the augmenting path $(z, y, x, u)$ in $\mathcal{G}$ (w.r.t. $M$ ) by
setting $M \leftarrow(M \backslash\{(x, y)\}) \cup\{(y, z),(x, u)\}$.
For all $w \in \operatorname{NeIGHBors}(u)$ :
Remove $u$ from the list Free-Neighbors $(w)$.
For all $w \in \operatorname{Neighbors}(z)$ :
Remove $z$ from the list Free-Neighbors $(w)$.
RETURN.
Figure 13: Resolve $(u)$.
If we are in Case 1, then the procedure runs for $O(1)$ time. In contrast, if we are in Case 2, then the procedure runs for $O(\Delta)$ time as both Find-mate and Resolve take time $O(\Delta)$. So the overall update
time for handling an edge deletion is $O(\Delta)$.

### 3.2 The kernel and its properties.

In the input graph $G=(V, E)$, let $\mathcal{N}_{v}=\{u \in V:(u, v) \in E\}$ denote the set of neighbors of $v \in V$.
Consider a subgraph $\kappa(G)=(V, \kappa(E))$ with $\kappa(E) \subseteq E$. For all $v \in V$, define the set $\kappa\left(\mathcal{N}_{v}\right)=\{u \in$ $\left.\mathcal{N}_{v}:(u, v) \in \kappa(E)\right\}$. If $u \in \kappa\left(\mathcal{N}_{v}\right)$, then we say that $u$ is a friend of $v$ in $\kappa(G)$. Next, the set of nodes $V$ is partitioned into two groups: tight and slack. We denote the set of tight (resp. slack) nodes by $\kappa_{T}(V)$ (resp. $\kappa_{S}(V)$ ). Thus, we have $V=\kappa_{T}(V) \cup \kappa_{S}(V)$ and $\kappa_{T}(V) \cap \kappa_{S}(V)=\emptyset$.

Definition 3.5. Fix any $c \geq 1$ and any $\epsilon \in[0,1 / 3)$. The subgraph $\kappa(G)$ is an $(\epsilon, c)$-kernel of $G$ with respect to the partition $\left(\kappa_{T}(V), \kappa_{S}(V)\right)$ iff it satisfies Invariants 3.6-3.8.

Invariant 3.6. $\left|\kappa\left(\mathcal{N}_{v}\right)\right| \leq(1+\epsilon)$ c for all $v \in V$, i.e., every node has at most $(1+\epsilon)$ c friends.
Invariant 3.7. $\left|\kappa\left(\mathcal{N}_{v}\right)\right| \geq(1-\epsilon) c$ for all $v \in \kappa_{T}(V)$, i.e., every tight node has at least $(1-\epsilon) c$ friends.
Invariant 3.8. For all $u, v \in \kappa_{S}(V)$, if $(u, v) \in E$, then $(u, v) \in \kappa(E)$. In other words, if two slack nodes are connected by an edge in $G$, then that edge must belong to $\kappa(G)$.

By Invariant 3.6, the maximum degree of a node in an $(\epsilon, c)$-kernel is $O(c)$. In Theorems 3.9 and 3.10, we show that an $(\epsilon, c)$-kernel $\kappa(G)$ is basically a subgraph of $G$ with maximum degree $O(c)$ that approximately preserves the size of the maximum matching.

For ease of exposition, we often refer to a kernel $\kappa(G)$ without explicitly mentioning the underlying partition $\left(\kappa_{T}(V), \kappa_{S}(V)\right)$. The proofs of Theorems 3.9, 3.10 appear in Sections 3.2.1 and 3.2.2 respectively.

Theorem 3.9. Let $M$ be a maximal matching in an $(\epsilon, c)$-kernel $\kappa(G)$. Then $M$ is a $(4+6 \epsilon)$-approximation to the maximum matching in $G$.

Theorem 3.10. Let $M$ be a matching in an $(\epsilon, c)$-kernel $\kappa(G)$ such that every augmenting path in $\kappa(G)$ (w.r.t. $M$ ) has length at least five. Then $M$ is a $(3+3 \epsilon)$-approximation to the maximum matching in $G$.

The next theorem shows that the approximation ratio derived in Theorem 3.9 is essentially tight. Its proof appears in Section 3.2.3.

Theorem 3.11. For every positive even integer $c$, there is a graph $G=(V, E), a(1 / c, c)$-kernel $\kappa(G)=$ $(V, \kappa(E))$, and a maximal matching $M$ in $\kappa(G)$ such that the size of the maximum matching in $G$ is at least 4 times the size of $M$.

### 3.2.1 Proof of Theorem 3.9.

Let $V(M)=\left\{v \in V:(u, v) \in M\right.$ for some $\left.u \in \mathcal{N}_{v}\right\}$ be the set of nodes that are matched in $M$, and let $F_{T}=\kappa_{T}(V) \backslash V(M)$ be the subset of tight nodes in $\kappa(G)$ that are free in $M$.

Lemma 3.12. Recall that $M$ is a maximal matching in the $(\epsilon, c)$-kernel $\kappa(G)$. We have $\left|F_{T}\right| \leq(2+6 \epsilon) \cdot|M|$.
Proof. We will show that $\left|F_{T}\right| \leq(1+3 \epsilon) \cdot|V(M)|$. Since $|V(M)|=2 \cdot|M|$, the lemma follows.
We design a charging scheme where each node in $F_{T}$ contributes one dollar to a global fund. So the total amount of money in this fund is equal to $\left|F_{T}\right|$ dollars. Below, we demonstrate how to transfer this fund to the nodes in $V(M)$ so that each node $x \in V(M)$ receives at most $(1+3 \epsilon)$ dollars. This implies that $\left|F_{T}\right| \leq(1+3 \epsilon) \cdot|V(M)|$.

Since $M$ is a maximal matching in $\kappa(G)$, we must have $\kappa\left(\mathcal{N}_{v}\right) \subseteq V(M)$ for all $v \in F_{T}$. For each node $v \in F_{T}$, we distribute its one dollar equally among its friends, i.e., each node $x \in \kappa\left(\mathcal{N}_{v}\right)$ gets $1 /\left|\kappa\left(\mathcal{N}_{v}\right)\right|$
dollars from $v$. Since $F_{T} \subseteq \kappa_{T}(V)$, Invariant 3.7 implies that $1 /\left|\kappa\left(\mathcal{N}_{v}\right)\right| \leq 1 /((1-\epsilon) c)$ for all $v \in F_{T}$. In other words, a node in $V(M)$ receives at most $1 /((1-\epsilon) c)$ dollars from each of its friends under this money-transfer scheme. But, by Invariant 3.6, a node can have at most $(1+\epsilon) c$ friends. So the total amount of money received by a node in $V(M)$ is at $\operatorname{most}(1+\epsilon) c /((1-\epsilon) c) \leq(1+3 \epsilon)$, for $\epsilon \in[0,1 / 3)$.

To continue with the proof of Theorem 3.9 , let $M^{o} \subseteq E$ be a maximum-cardinality matching in $G=$ $(V, E)$. Define $M_{1}^{o} \subseteq M^{o}$ to be the subset of edges whose both endpoints are unmatched in $M$, and let $M_{2}^{o}=M^{o} \backslash M_{1}^{o}$.

Consider any edge $(u, v) \in M_{1}^{o}$. By definition, both the nodes $u, v$ are free in $M$. Since $M$ is a maximal matching in $\kappa(G)$, it follows that the edge $(u, v)$ is not part of the kernel, i.e., $(u, v) \neq \kappa(E)$. Hence, Invariant 3.8 implies that either $u \notin \kappa_{S}(V)$ or $v \notin \kappa_{S}(V)$. Without any loss of generality, suppose that $u \notin \kappa_{S}(V)$. This means that the node $u$ is tight, and, furthermore, it is free in $M$. We infer that every edge in $M_{1}^{o}$ is incident to at least one node from $F_{T}$, where $F_{T} \subseteq \kappa_{T}(V)$ is the subset of tight nodes that are free in $M$. Accordingly, we have $\left|M_{1}^{o}\right| \leq\left|F_{T}\right|$. Combining this inequality with Lemma 3.12, we get:

$$
\begin{equation*}
\left|M_{1}^{o}\right| \leq(2+6 \epsilon) \cdot|M| \tag{23}
\end{equation*}
$$

Next, every edge in $M_{2}^{o}$ has at least one endpoint that is matched in $M$. Thus, $M$ is a maximal matching in the graph $G^{\prime}=\left(V, M_{2}^{o} \cup M\right)$. Since $M_{2}^{o}$ is also a matching in $G^{\prime}$, we get:

$$
\begin{equation*}
\left|M_{2}^{o}\right| \leq 2 \cdot|M| \tag{24}
\end{equation*}
$$

The theorem follows if we add equations 23 and 24.

### 3.2.2 Proof of Theorem 3.10.

Let $V(M)=\left\{v \in V:(u, v) \in M\right.$ for some $\left.u \in \mathcal{N}_{v}\right\}$ be the set of nodes that are matched in $M$. For all nodes $v \in V(M)$, let $e_{M}(v)$ denote the edge in $M$ that is incident to $v$. Let $F_{T}=\kappa_{T}(V) \backslash V(M)$ be the subset of tight nodes in $\kappa(G)$ that are free in $M$.
Lemma 3.13. Recall that $M$ is a matching in the $(\epsilon, c)$-kernel $\kappa(G)$ such that every augmenting path in $\kappa(G)$ (w.r.t. $M$ ) has length at least five. For every edge $(u, v) \in M$, we have $\mid\left(\kappa\left(\mathcal{N}_{u}\right) \cap F_{T}\right) \cup$ $\left(\kappa\left(\mathcal{N}_{v}\right) \cap F_{T}\right) \mid \leq(1+\epsilon) c$.

Proof. Suppose that the lemma is false and we have $\left|\left(\kappa\left(\mathcal{N}_{u}\right) \cap F_{T}\right) \cup\left(\kappa\left(\mathcal{N}_{v}\right) \cap F_{T}\right)\right|>(1+\epsilon) c$ for some edge $(u, v) \in M$. As Invariant 3.6 guarantees that $\left|\kappa\left(\mathcal{N}_{u}\right)\right| \leq(1+\epsilon) c$ and $\left|\kappa\left(\mathcal{N}_{v}\right)\right| \leq(1+\epsilon) c$, there has to be a pair of distinct nodes $u^{\prime}, v^{\prime} \in F_{T}$ such that $u^{\prime} \in \kappa\left(\mathcal{N}_{u}\right)$ and $v^{\prime} \in \kappa\left(\mathcal{N}_{v}\right)$. This means that the path $\left(u^{\prime}, u, v, v^{\prime}\right)$ is an augmenting path in $\kappa(G)$ (w.r.t. $\left.M\right)$ and has length three. We reach a contradiction.

Lemma 3.14. Recall that $M$ is a matching in the $(\epsilon, c)$-kernel $\kappa(G)$ such that every augmenting path in $\kappa(G)\left(\right.$ w.r.t. $M$ ) has length at least five. We have $\left|F_{T}\right| \leq(1+3 \epsilon) \cdot|M|$.
Proof. We design a charging scheme where each node in $F_{T}$ contributes one dollar to a global fund. So the total amount of money in this fund is equal to $\left|F_{T}\right|$ dollars. Below, we demonstrate how to transfer this fund to the edges in $M$ so that each edge $e \in M$ receives at most $(1+3 \epsilon)$ dollars. This implies that $\left|F_{T}\right| \leq(1+3 \epsilon) \cdot|M|$.

Since $M$ is a maximal matching in $\kappa(G)$, we must have $\kappa\left(\mathcal{N}_{v}\right) \subseteq V(M)$ for all $v \in F_{T}$. For each node $v \in F_{T}$, we distribute its one dollar equally among the matched edges incident to its friends, i.e., for each node $x \in \kappa\left(\mathcal{N}_{v}\right)$, the edge $e_{M}(x)$ gets $1 /\left|\kappa\left(\mathcal{N}_{v}\right)\right|$ dollars from $v$. Since $F_{T} \subseteq \kappa_{T}(V)$, Invariant 3.7 implies that $1 /\left|\kappa\left(\mathcal{N}_{v}\right)\right| \leq 1 /((1-\epsilon) c)$ for all $v \in F_{T}$. In other words, an edge $(x, y) \in M$ receives at most $1 /((1-\epsilon) c)$ dollars from each of the nodes $v \in\left(\kappa\left(\mathcal{N}_{x}\right) \cap F_{T}\right) \cup\left(\kappa\left(\mathcal{N}_{y}\right) \cap F_{T}\right)$ under this money-transfer scheme. Hence, Lemma 3.13 implies that the total amount of money received by an edge $(x, y) \in M$ is at $\operatorname{most}(1+\epsilon) c /((1-\epsilon) c) \leq(1+3 \epsilon)$, for $\epsilon \in[0,1 / 3)$.

To continue with the proof of Theorem 3.10, define $M^{o} \subseteq E$ to be a maximum-cardinality matching in $G=(V, E)$. Furthermore, let $M_{1}^{o} \subseteq M^{o}$ be the subset of edges whose both endpoints are unmatched in $M$, and let $M_{2}^{o}=M^{o} \backslash M_{1}^{o}$.

Consider any edge $(u, v) \in M_{1}^{o}$. By definition, both the nodes $u, v$ are free in $M$. Since $M$ is a maximal matching in $\kappa(G)$, it follows that the edge $(u, v)$ is not part of the kernel, i.e., $(u, v) \neq \kappa(E)$. Hence, Invariant 3.8 implies that either $u \notin \kappa_{S}(V)$ or $v \notin \kappa_{S}(V)$. Without any loss of generality, suppose that $u \notin \kappa_{S}(V)$. This means that the node $u$ is tight, and, furthermore, it is free in $M$. We infer that every edge in $M_{1}^{o}$ is incident to at least one node from $F_{T}$. Accordingly, we have $\left|M_{1}^{o}\right| \leq\left|F_{T}\right|$. Combining this inequality with Lemma 3.14, we get:

$$
\begin{equation*}
\left|M_{1}^{o}\right| \leq(1+3 \epsilon) \cdot|M| \tag{25}
\end{equation*}
$$

Next, every edge in $M_{2}^{o}$ has at least one endpoint that is matched in $M$. Thus, $M$ is a maximal matching in the graph $G^{\prime}=\left(V, M_{2}^{o} \cup M\right)$. Since $M_{2}^{o}$ is also a matching in $G^{\prime}$, we get:

$$
\begin{equation*}
\left|M_{2}^{o}\right| \leq 2 \cdot|M| \tag{26}
\end{equation*}
$$

The theorem follows if we add equations 25 and 26.

### 3.2.3 Proof of Theorem 3.11

We define the set of nodes $U=\left\{u_{1}, \ldots, u_{c / 2}\right\}, U^{\prime}=\left\{u_{1}^{\prime}, \ldots, u_{c / 2}^{\prime}\right\}, U^{*}=\left\{u_{1}^{*}, \ldots, u_{c}^{*}\right\}, X=\left\{x_{1}, \ldots, x_{c}\right\}$, $Y=\left\{y_{1}, \ldots, y_{c / 2}\right\}$ and $Z=\left\{z_{1}, \ldots, z_{c / 2}\right\}$. Next, we define the following sets of edges.

$$
\begin{aligned}
E\left(U, U^{\prime}\right) & =\bigcup_{i=1}^{c / 2}\left\{\left(u_{i}, u_{i}^{\prime}\right)\right\} \\
E\left(U, U^{*}\right) & =\bigcup_{u \in U, u^{*} \in U^{*}}\left\{\left(u, u^{*}\right)\right\} \\
E\left(U^{\prime}, U^{*}\right) & =\bigcup_{u^{\prime} \in U, u^{*} \in U^{*}}\left\{\left(u^{\prime}, u^{*}\right)\right\} \\
E\left(X, U^{*}\right) & =\bigcup_{i=1}^{c}\left\{\left(x_{i}, u_{i}^{*}\right)\right\} \\
E\left(Y, U, U^{\prime}, Z\right) & =\bigcup_{i=1}^{c / 2}\left\{\left(y_{i}, u_{i}\right),\left(u_{i}^{\prime}, z_{i}\right)\right\}
\end{aligned}
$$

Next, we consider a graph $G=(V, E)$ defined on the node-set $V=U \cup U^{\prime} \cup U^{*} \cup X \cup Y \cup Z$ and edge-set $E=E\left(U, U^{\prime}\right) \cup E\left(U, U^{*}\right) \cup E\left(U^{\prime}, U^{*}\right) \cup E\left(X, U^{*}\right) \cup E\left(Y, U, U^{\prime}, Z\right)$. See Figure 14. Define the subset of edges $\kappa(E)=E\left(U, U^{\prime}\right) \cup E\left(U, U^{*}\right) \cup E\left(U^{\prime}, U^{*}\right)$. It is easy to check that the subgraph $\kappa(G)=(V, \kappa(E))$ is a $(1 / c, c)$-kernel in $G$, with the set of tight nodes being $\kappa_{T}(V)=U \cup U^{\prime} \cup U^{*}$, and the set of slack nodes being $\kappa_{S}(V)=X \cup Y \cup Z$. We now consider two matchings $M$ and $M^{*}$ defines as follows.

$$
\begin{aligned}
M & =\bigcup_{i=1}^{c / 2}\left\{\left(u_{i}, u_{i}^{\prime}\right)\right\} \\
M^{*} & =E\left(X, U^{*}\right) \bigcup E\left(Y, U, U^{\prime}, Z\right)
\end{aligned}
$$

The theorem follows from the observations that $M$ is a maximal matching in $\kappa(G)$ and $|M|=c / 2,\left|M^{*}\right|=$ $2 c$.


Figure 14: The graph considered in the proof of Theorem 3.11, for $c=4$.

### 3.3 An algorithm for building a kernel in a static graph.

We present a linear-time algorithm for constructing a $(0, c)$-kernel $\kappa(G)$ of a static graph $G=(V, E)$.
Theorem 3.15. We have an algorithm for computing a $(0, c)$-kernel $\kappa(G)=(V, \kappa(E))$ of a graph $G=$ $(V, E)$. The kernel returned by the algorithm has the added property that tight nodes have exactly $c$ friends and slack nodes have less than c friends. For every $c \geq 1$, the algorithm runs in $O(|E|)$ time.

Proof. Initially, each node in $V$ has zero friends, and the edge-set $\kappa(E)$ is empty. We then execute the For loop stated below.

- For all $(u, v) \in E$
- IF $\left|\kappa\left(\mathcal{N}_{u}\right)\right|<c$ and $\left|\kappa\left(\mathcal{N}_{v}\right)\right|<c$, THEN
* Set $\kappa\left(\mathcal{N}_{v}\right) \leftarrow \kappa\left(\mathcal{N}_{v}\right) \cup\{u\}$ and $\kappa\left(\mathcal{N}_{u}\right) \leftarrow \kappa\left(\mathcal{N}_{u}\right) \cup\{v\}$.
* Set $\kappa(E) \leftarrow \kappa(E) \cup\{(u, v)\}$.

Consider the $\kappa(G)=(V, \kappa(E))$ we get at the end of the For loop. Clearly, in $\kappa(G)$ every node has at most $c$ friends, i.e., $\left|\kappa\left(\mathcal{N}_{v}\right)\right| \leq c$ for all $v \in V$. Furthermore, for all edges $(u, v) \in E$, if $\left|\kappa\left(\mathcal{N}_{u}\right)\right|<c$ and $\left|\kappa\left(\mathcal{N}_{v}\right)\right|<c$, then it is guaranteed that $(u, v) \in \kappa(E)$. Thus, if we define $\kappa_{T}(V)=\left\{v \in V:\left|\kappa\left(\mathcal{N}_{v}\right)\right|=c\right\}$ and $\kappa_{S}(V)=\left\{v \in V:\left|\kappa\left(\mathcal{N}_{v}\right)\right|<c\right\}$, then $\kappa(G)$ becomes a $(0, c)$-kernel of $G$. The whole procedure can be implemented in $O(|E|)$ time.

### 3.4 Data structures for representing a kernel in a dynamic graph.

In the graph $G$, the set of neighbors of a node $v$ is stored in the form of a linked list Neighbors $(v)$, which is part of an adjacency-list data structure. Further, each node $v$ maintains the following information.

- A bit Type $(v)$ indicating if the node is tight or slack.

$$
\operatorname{TYPE}(v)= \begin{cases}\text { tight } & \text { if } v \in \kappa_{T}(V) \\ \text { slack } & \text { if } v \in \kappa_{S}(V)\end{cases}
$$

- The set of nodes $\kappa\left(\mathcal{N}_{v}\right)$, in the form of a doubly linked list FriEnds $(v)$.
- A counter \#FRIENDS $(v)=\left|\kappa\left(\mathcal{N}_{v}\right)\right|$ that keeps track of the number of friends of the node.

Furthermore, we store at each edge $(u, v) \in \kappa(E)$ two pointers, corresponding to the two occurrences of edge $(u, v)$ in the linked lists $\operatorname{Friends}(u)$ and $\operatorname{Friends}(v)$. In particular, we denote by Pointer $[u, v]$ (resp. Pointer $[v, u]$ ) the pointer to the position of $u$ (resp. $v$ ) in the list Friends $(v)$ (resp. Friends $(u)$ ). Using those pointers, we can insert an edge into $\kappa(G)$ or delete an edge from $\kappa(G)$ in $O(1)$ time.

## $4(3+\epsilon)$-approximate matching in $O(\sqrt{n} / \epsilon)$ amortized update time

Fix any $\epsilon \in(0,1 / 3)$. We present an algorithm for maintaining a $(3+\epsilon)$-approximate matching in a graph $G=(V, E)$ undergoing a sequence of updates with $O(\sqrt{n} / \epsilon)$ amortized update time, where $n=|V|$ denotes the number of nodes in $G$.

### 4.1 Overview of our approach.

Before the graph $G=(V, E)$ starts changing, we preprocess it in linear time to build a $(0, \sqrt{n})$-kernel $\kappa(G)$ of $G$ (see Theorem 3.15). Next, we compute a matching $M$ in $\kappa(G)$ such that every augmenting path in $\kappa(G)$ (w.r.t. $M$ ) has length at least five. This also takes $O(|E|)$ time.

Thus, initially we have a $(0, \sqrt{n})$-kernel of $G$, which is also an $(\epsilon, \sqrt{n})$-kernel of $G$ for all $\epsilon \geq 0$. After each update (edge insertion/deletion) in $G$, we first modify $\kappa(G)$ using the algorithm in Section 4.2, and subsequently we modify the matching $M$ in $\kappa(G)$ as per Theorem 3.4. The algorithm in Section 4.2 guarantees that $\kappa(G)$ remains an $(\epsilon, \sqrt{n})$-kernel of $G$ after this update. Hence, by Theorem 3.10, the matching $M$ always remains a $(3+\epsilon)$-approximation to the maximum matching in $G$.

We now explain the idea behind the algorithm in Section 4.2. In the beginning of the phase, Theorem 3.15 guarantees that every tight node has exactly $\sqrt{n}$ friends. During the phase, whenever the number of friends of a tight node $v$ drops below $(1-\epsilon) \sqrt{n}$, it scans through its entire list of neighbors in $\mathcal{N}_{v}$. While passing through a node $u \in \mathcal{N}_{v}$ during this scan, if it discovers that both $\left|\kappa\left(\mathcal{N}_{u}\right)\right|,\left|\kappa\left(\mathcal{N}_{v}\right)\right|<\sqrt{n}$ and $u \notin \kappa\left(\mathcal{N}_{v}\right)$, then it adds the edge $(u, v)$ to $\kappa(G)$. This procedure is implemented in the subroutine REFILL-NOW-FRIENDS $(v)$ (see Figure 15). Note that this subroutine can take $O(n)$ time, but it is called only when a tight node has lost $\epsilon \sqrt{n}$ of its friends. Each such loss corresponds to an edge deletion in $\kappa(E) \subseteq E$. Thus, distributing this cost among the concerned edge deletions in $G$, we get an amortized update time of $O(\sqrt{n} / \epsilon)$.

Finally, in Section 4.3, we show that the amortized update time for maintaining the matching $M$ in $\kappa(G)$ is also $O(\sqrt{n} / \epsilon)$, and this leads to the main result of this section.

Theorem 4.1. In a dynamic graph $G=(V, E)$, we can maintain a $(3+\epsilon)$-approximate matching $M \subseteq E$ in $O(\sqrt{n} / \epsilon)$ amortized update time.

### 4.2 Maintaining the kernel $\kappa(G)$.

We will ensure that the kernel $\kappa(G)$ always satisfies the five properties described below. Clearly, these properties hold immediately after the preprocessing step (see Theorem 3.15).

- (P1) Every node has at most $\sqrt{n}$ friends in $\kappa(G)$.
- (P2) Every tight node has at least $(1-\epsilon) \sqrt{n}$ friends in $\kappa(G)$.
- (P3) If two slack nodes are connected by an edge in $G$, then they are friends of each other in $\kappa(G)$.
- (P4) If a node has exactly $\sqrt{n}$ friends in $\kappa(G)$, then it is tight.
- (P5) A slack node $v$ becomes tight only when $\left|\kappa\left(\mathcal{N}_{v}\right)\right|$ increases from $\sqrt{n}-1$ to $\sqrt{n}$.

We now describe our algorithm for updating $\kappa(G)$ after an edge insertion/deletion in $G$. It is easy to check that if (P1)-(P5) were true before the update, then they continue to be true after the update.

The three properties (P1)-(P3) imply that $\kappa(G)$ is an $(\epsilon, \sqrt{n})$-kernel. The remaining properties (P4)-(P5) are helpful for reasoning about the correctness of our algorithm and for bounding its amortized update time.

Handling an edge insertion in $G$. Suppose that an edge $(u, v)$ is inserted into the graph $G=(V, E)$. To handle this edge insertion, we first update the lists neighbors ( $u$ ) and Neighbors $(v)$, and then process the edge as follows.

- Case 1: $\operatorname{Either} \# \operatorname{Friends}(u)=\sqrt{n}$ or $\# \operatorname{Friends}(v)=\sqrt{n}$.

We do nothing.

- Case 2: Both \#Friends $(u)<\sqrt{n}$ and \#Friends $(v)<\sqrt{n}$.

We add the edge $(u, v)$ to the kernel $\kappa(G)$ and make $u, v$ friends of each other. Specifically, we add the node $u$ to the set Friends $(v)$, the node $v$ to the set Friends $(u)$, update the pointers Pointer $[u, v]$, Pointer $[v, u]$, and increment each of the counters \#Friends $(u)$, \#Friends $(v)$ by one unit. After this increment, if \#Friends $(u)=\sqrt{n}$, then we set $\operatorname{TyPE}(u) \leftarrow$ tight. Similarly, if \#Friends $(v)=$ $\sqrt{n}$, then we set $\operatorname{Type}(v) \leftarrow$ tight.

Handling an edge deletion in $G$. Suppose that an edge $(u, v)$ gets deleted from the graph $G=(V, E)$. To handle this edge deletion, we first update the lists NEighbors $(u)$ and NEighbors $(v)$.

Next, we check if the edge $(u, v)$ was part of the kernel $\kappa(G)$, and, if yes, then we delete $(u, v)$ from $\kappa(E)$. Specifically, we delete $u$ from $\operatorname{Friends}(v), v$ from $\operatorname{Friends}(u)$, update the pointers Pointer $[u, v]$, Pointer $[v, u]$, and decrement each of the counters \#Friends $(u)$, \#Friends $(v)$ by one unit.

We then process the nodes $u$ and $v$ one after another. Below, we describe only the procedure that runs on the node $u$. The procedure for the node $v$ is exactly the same.

- Case 1: $\operatorname{Type}(u)=$ tight. Here, we check if the number of friends of $u$ has dropped below the prescribed limit due to the edge deletion, and, accordingly, we consider two possible subcases.
- Case 1a: \#Friends $(u)<(1-\epsilon) \sqrt{n}$. Here, we call the subroutine REFILL-NOW-FRIENDS $(u)$ described in Figure 15. This takes $O(n)$ time.
- Case 1b: \#Friends $(u) \geq(1-\epsilon) \sqrt{n}$. We do nothing in this case.
- Case 2: $\operatorname{Type}(u)=$ slack. We do nothing in this case.

Bounding the amortized update time. If we ignore the calls to the subroutine REFILL-NOW-FRIENDS(.), then handling every edge insertion/deletion in $G$ requires $O(1)$ time in the worst case. We will show that the time taken by all the calls to REFILL-NOW-FRIENDS(.), amortized over a sequence of updates in $G$, is $O(\sqrt{n} / \epsilon)$.

Fix any sequence of $t$ edge updates in $G$, and let $t_{v}$ denote the number of edge updates in this sequence that are incident to a node $v \in V$. Since each edge is incident to two nodes, we have $2 t=\sum_{v \in V} t_{v}$. Let $r_{v}$ denote the number of times the subroutine REFILL-NOW-FRIENDS $(v)$ is called during this sequence, so that the total number of calls to the subroutine REFILL-NOW-FRIENDS(.) is given by $\sum_{v \in V} r_{v}$. Since

```
IF the list NEIGHBORS \((u)\) is empty, THEN
    Set Type \((u) \leftarrow\) slack.
    RETURN.
Let \(x\) be the first node in the list NEIGHBORS \((u)\).
While (\#FRIENDS \((u)<\sqrt{n}\) )
    IF \(x \notin \operatorname{FriEndS}(u)\) and \#Friends \((x)<\sqrt{n}\), Then
        Insert \(x\) into the list \(\operatorname{Friends}(u)\), and \(u\) into the list Friends \((x)\), so that the edge \((u, x)\)
        becomes part of the kernel \(\kappa(G)=(V, \kappa(E))\). Update the pointers Pointer \([u, x]\) and
        Pointer \([x, u]\). Increment each of the counters \#FRIENDS \((u)\), \#FRIENDS \((v)\) by one unit.
            After this increment, if \#Friends \((x)=\sqrt{n}\), then set \(\operatorname{TyPE}(x) \leftarrow\) tight.
    IF \(x\) is the last node in the list NEIGHBORS \((u)\), THEN
        exit the While loop.
    ELSE
            Let \(y\) be the node that succeeds \(x\) in the list NEIGHBORS \((u)\).
            Set \(x \leftarrow y\).
If \#FRIENDS \((u)<\sqrt{n}\), THEN
    Set TyPE \((u) \leftarrow\) slack.
```

Figure 15: REFILL-NOW-FRIENDS( $u$ ).
each individual call to REFILL-NOW-FRIENDS(.) takes $O(n)$ time in the worst case, it suffices to show the following guarantee.

$$
\begin{equation*}
\frac{\sum_{v} r_{v}}{t} \leq \frac{\alpha}{\epsilon \sqrt{n}}, \text { for some constant } \alpha>0 . \tag{27}
\end{equation*}
$$

Consider any node $v \in V$. Our preprocessing step guarantees that in the beginning of the algorithm, the node $v$ is tight only if it has exactly $\sqrt{n}$ friends in $\kappa(G)$. By (P5), during the sequence of updates in $G$, the node $v$ changes its status from slack to tight only if $\left|\kappa\left(\mathcal{N}_{v}\right)\right|$ becomes equal to $\sqrt{n}$. On the other hand, the subroutine REFILL-NOW-FRIENDS $(v)$ is called only if $v$ is tight and $\left|\kappa\left(\mathcal{N}_{v}\right)\right|$ drops below $(1-\epsilon) \sqrt{n}$. Accordingly, a call to REFILL-NOW-FRIENDS $(v)$ occurs when $v$ has lost at least $\epsilon \sqrt{n}$ friends in $\kappa(G)$, and each such loss stems from the deletion of an edge in $\kappa(E) \subseteq E$ that was incident to $v$. Thus, we conclude that $r_{v} / t_{v} \leq 1 /(\epsilon \sqrt{n})$. We now derive equation 27 as follows.

$$
\frac{\sum_{v} r_{v}}{t}=2 \cdot \frac{\sum_{v} r_{v}}{\sum_{v} t_{v}} \leq 2 \cdot \max _{v}\left(\frac{r_{v}}{t_{v}}\right)=\frac{2}{\epsilon \sqrt{n}} .
$$

Lemma 4.2. We can maintain an $(\epsilon, c)$-kernel $\kappa(G)$ of a dynamic graph $G$. Updating the kernel requires $O(\sqrt{n} / \epsilon)$ time, amortized over the number of edge insertions/deletions in $G$.

### 4.3 Maintaining the matching $M$ in $\kappa(G)$.

We maintain the matching $M$ in $\kappa(G)$ using Theorem 3.4. To bound the amortized update time of this procedure, fix any sequence of $t$ edge updates in $G$. Ignoring the calls to REFILL-NOW-FRIENDS(.), each update in $G$ leads to at most one update in $\kappa(G)$. On the other hand, every single call to REFILL-NOW-FRIENDS(.) leads to $O(\sqrt{n})$ updates in $\kappa(G)$. By equation 27, a sequence of $t$ updates in $G$ results in at most $O(t /(\epsilon \sqrt{n}))$ calls to REFILL-NOW-FRIENDS(.). Hence, such a sequence can cause at most $O(t / \epsilon)$ updates in $\kappa(G)$. As the maximum degree of a node in $\kappa(G)$ is $O(c)$, each update in $\kappa(G)$ is handled in $O(\sqrt{n})$ time by the algorithm in Theorem 3.4. This gives us a total running time of $O((t / \epsilon) \cdot \sqrt{n})$ for maintaining the matching $M$ during a sequence of $t$ updates in $G$. Thus, the amortized update time becomes $O(\sqrt{n} / \epsilon)$.

Lemma 4.3. In a dynamic graph $G$, we can maintain an $(\epsilon, c)$-kernel $\kappa(G)$ and a matching $M$ in $\kappa(G)$ such that every augmenting path in $\kappa(G)$ (w.r.t. $M$ ) has length at least five. Updating the kernel $\kappa(G)$ and the matching $M$ requires $O(\sqrt{n} / \epsilon)$ time, amortized over the number of egde insertions/deletions in $G$.

Proof of Theorem 4.1. By Theorem 3.10, the matching maintained by Lemma 4.3 is a $(3+\epsilon)$-approximate maximum matching in the input graph $G=(V, E)$.

## $5(3+\epsilon)$-approximate matching in $O\left(m^{1 / 3} / \epsilon^{2}\right)$ amortized update time

Fix any constant $\epsilon \in(0,1 / 3)$. In this section, we present an algorithm for maintaining a ( $3+\epsilon$ )-approximate matching $M$ in a graph $G=(V, E)$ undergoing a sequence of edge insertions/deletions. It requires $O\left(m^{1 / 3} / \epsilon^{2}\right)$ amortized update time.

### 5.1 Overview of our approach.

The main idea is to partition the sequence of updates (edge insertions/deletions) in $G$ into phases. Each phase lasts for $\epsilon^{2} c^{2} / 2$ consecutive updates in $G$. Let $G_{i, t}$ denote the state of $G$ just after the $t^{\text {th }}$ update in phase $i$, for any $0 \leq t \leq \epsilon^{2} c^{2} / 2$. The initial state of the graph, before it starts changing, is given by $G_{1,0}$. We reach the graph $G_{i, t}$ from $G_{1,0}$ after a sequence of $(i-1) \cdot\left(\epsilon^{2} c^{2} / 2\right)+t$ updates in $G$.

For the rest of this section, we focus on describing our algorithm for any given phase $i \geq 1$. We define $m \leftarrow\left|E_{i, 0}\right|$ to be the number of edges in the input graph in the beginning of the phase, and set $c=m^{1 / 3}$. Since the phase lasts for only $O\left(\epsilon^{2} m^{2 / 3}\right)$ updates in $G$, it follows that $\left|E_{i, t}\right|=O(m)$ for all $0 \leq t \leq \epsilon^{2} c^{2} / 2$. During the phase, we maintain a $(3+\epsilon)$-approximate matching in $G$ as described below.

Just before phase $i$ begins, we build a $(0, c)$-kernel $\kappa(G)$ on $G=G_{i, 0}$ as per Theorem 3.15. This takes $O(m)$ time. Next, we compute a $3 / 2$-approximate maximum matching $M=M_{i, 0}$ on $\kappa\left(G_{i, 0}\right)$ by ensuring that every augmenting path in $\kappa\left(G_{i, 0}\right)$ (w.r.t. $M_{i, 0}$ ) has length at least five. This also takes $O(m)$ time, and concludes the preprocessing step.

After each update in $G$ during phase $i$, we first modify the kernel $\kappa(G)$ using the algorithm in Section 5.2, and subsequently we modify the matching $M$ in $\kappa(G)$ using Theorem 3.4. Theorem 5.2 below guarantees that the graph $\kappa(G)$ remains an $(\epsilon, c)$-kernel of $G$ throughout the phase. Hence, by Theorem 3.10, the matching $M$ remains a $(3+\epsilon)$-approximation to the maximum matching in $G$.

The idea behind the algorithm in Section 5.2 is simple. In the beginning of the phase, Theorem 3.15 guarantees that every tight node has exactly $c$ friends. During the phase, whenever the number of friends of a tight node $v$ drops below $(1-\epsilon) c$, it scans through its first $c$ neighbors in $\mathcal{N}_{v}$, and keeps making friends out of them until $\left|\kappa\left(\mathcal{N}_{v}\right)\right|$ becomes equal to $c$. This procedure is implemented in the subroutine REFILL $(v)$ (see Figure 16). A potential problem with this approach is that after a while a node $v$ may have more than $(1+\epsilon) c$ friends due to the repeated invocations of the subroutine $\operatorname{REFILL}(u)$, for neighbors $u \in \mathcal{N}_{v}$. We show that this event can be ruled out by ending the current phase after $\epsilon^{2} c^{2} / 2$ updates in $G$.

## Bounding the update time.

- Preprocessing.

The preprocessing in the beginning of the phase takes $O(m)$ time. Since the phase lasts for $\epsilon^{2} c^{2} / 2$ updates, we get an amortized bound of $O\left(m /\left(\epsilon^{2} c^{2}\right)\right)=O\left(m^{1 / 3} / \epsilon^{2}\right)$.

- Maintaining the kernel $\kappa(G)$.

By Theorem 5.2 below, the kernel $\kappa(G)$ can be modified after each update in $G$ in $O(c)=O\left(m^{1 / 3}\right)$ time.

- Maintaining the matching $M$ in $\kappa(G)$.

By Theorem 5.2 below, the number of updates made into $\kappa(G)$ during the entire phase is $O\left(\epsilon^{2} c^{2}\right)$. Since the maximum degree of a node in $\kappa(G)$ is $O(c)$, modifying the matching $M$ after each update in $\kappa(G)$ requires $O(c)$ time (see Theorem 3.4). Thus, the total time spent during the phase in maintaining the matching $M$ is $O\left(\epsilon^{2} c^{2}\right) \cdot O(c)=O\left(\epsilon^{2} c^{3}\right)$. Since the phase lasts for $\epsilon^{2} c^{2} / 2$ updates in $G$, we get an amortized bound of $O(c)=O\left(m^{1 / 3}\right)$.
We summarize the main result of Section 5 in the theorem below. Note that in order to prove Theorem 5.1, we only need to show how to maintain an $(\epsilon, c)$-kernel in $O(c)$ update time. We do this in Section 5.2.
Theorem 5.1. In a dynamic graph $G=(V, E)$, we can maintain a $(3+\epsilon)$-approximate matching $M \subseteq E$ in $O\left(\mathrm{~m}^{1 / 3} / \epsilon^{2}\right)$ amortized update time.

### 5.2 Algorithms for maintaining the kernel during a phase.

We present our algorithm for maintaining an $(\epsilon, c)$-kernel of the graph during a phase.
Theorem 5.2. Suppose that we are given a $(0, c)$-kernel $\kappa(G)$ of $G$ in the beginning of a phase. Then we have an algorithm for modifying $\kappa(G)$ after each update (edge insertion/deletion) in $G$ such that:

1. Each update in $G$ is handled in $O(c)$ worst-case time.
2. At most $O\left(\epsilon^{2} c^{2}\right)$ updates are made into $\kappa(G)$ during the entire phase.
3. Throughout the duration of the phase, $\kappa(G)$ remains an $(\epsilon, c)$-kernel of $G$.

The rest of Section 5.2 is organized as follows. In Section 5.2.1 (resp. 5.2.2), we present our algorithm for handling an edge insertion (resp. edge deletion) in $G$ during the phase. In Section 5.2.3, we prove that the algorithm satisfies the properties stated in Theorem 5.2.

### 5.2.1 Handling an edge insertion in $G$ during the phase.

Suppose that an edge $(u, v)$ is inserted into the graph $G=(V, E)$. To handle this edge insertion, we first update the lists NEighbors $(u)$ and Neighbors $(v)$, and then process the edge as follows.

- Case 1: $\operatorname{Either} \operatorname{Type}(u)=$ tight or $\operatorname{Type}(v)=$ tight.

We do nothing and conclude the procedure.

- Case 2: Both $\operatorname{Type}(u)=$ slack and $\operatorname{Type}(v)=$ slack.
- Case 2a: Either \#Friends $(u) \geq c$ or \#Friends $(v) \geq c$.

If \#Friends $(u) \geq c$, then we set $\operatorname{TyPE}(u) \leftarrow$ tight. Next, if \#Friends $(v) \geq c$, then we set $\operatorname{Type}(v) \leftarrow$ tight.

- Case 2b: Both \#Friends $(u)<c$ and \#Friends $(v)<c$.

We add the edge $(u, v)$ to the kernel $\kappa(G)$ and make $u, v$ friends of each other. Specifically, we add the node $u$ to the list $\operatorname{Friends}(v)$ and $v$ to the list $\operatorname{Friends}(u)$, update the pointers Pointer $[u, v]$ and Pointer $[v, u]$ accordingly, and increment each of the counters \#Friends $(u)$, \#Friends $(v)$ by one unit.

Lemma 5.3. Suppose that an edge is inserted into the graph $G$ and we run the procedure described above.

1. This causes at most one edge insertion into $\kappa(G)$ and at most one edge deletion from $\kappa(G)$.
2. The procedure runs in $O(1)$ time in the worst case.

### 5.2.2 Handling an edge deletion in $G$ during the phase.

Suppose that an edge $(u, v)$ is deleted from the graph $G=(V, E)$. To handle this edge deletion, we first update the lists NEighbors $(u)$ and NEighbors $(v)$, and then proceed as follows.

We check if the edge $(u, v)$ was part of the kernel $\kappa(G)$, and, if yes, then we delete $(u, v)$ from $\kappa(G)$. Specifically, we delete $u$ from $\operatorname{Friends}(v), v$ from Friends $(u)$ (using Pointer $[u, v]$ and Pointer $[v, u]$ ), and decrement each of the counters \#Friends $(u)$, \#Friends $(v)$ by one unit.

We then process the nodes $u$ and $v$ one after another. Below, we describe only the procedure that runs on the node $u$. The procedure for the node $v$ is exactly the same.

- Case 1: $\operatorname{Type}(u)=$ tight. Here, we check if the number of friends of $u$ has dropped below the prescribed limit due to the edge deletion, and, accordingly, we consider two possible sub-cases.
- Case 1a: \#Friends $(u)<(1-\epsilon) c$. Here, we call the subroutine $\operatorname{REFILL}(u)$ as described in Figure 16.
- Case 1b: \#Friends $(u) \geq(1-\epsilon) c$. In this case, we do nothing and conclude the procedure.
- Case 2: $\operatorname{Type}(u)=$ slack. In this case, we do nothing and conclude the procedure.

Remark. Note the main difference between the REFILL-NOW-FRIENDS(u) procedure in Section 4 (see Figure 15) and the $\operatorname{REFILL}(u)$ procedure in this section (see Figure 16). The former procedure (Figure 15) ensures that no node ever gets more than $\sqrt{n}$ friends. On the other hand, in this section a node $x \notin$ Friends $[u]$ can become a friend of $u$ no matter how many friends $x$ already has. Thus, a node $x$ can possibly get more than $c$ friends. A phase, however, lasts for no more than $\epsilon^{2} c^{2} / 2$ edge insertion/deletions. Using this fact, we will show that a node $x$ can have at most $(1+\epsilon) c$-many friends at any point in time.

Lemma 5.4. Suppose that an edge is deleted from the graph $G$ and we run the procedure described above.

1. This causes at most $O(\epsilon c)$ edge insertions into $\kappa(G)$ and at most one edge deletion from $\kappa(G)$.
2. The procedure runs in $O(c)$ time in the worst case.

### 5.2.3 Proof of Theorem 5.2.

The first part of the theorem immediately follows from Lemmas 5.3, 5.4. We now focus on the second part. Note that at most one edge is deleted from $\kappa(G)$ after an update (edge insertion or deletion) in $G$. Since the phase lasts for $\epsilon^{2} c^{2} / 2$ updates in $G$, at most $\epsilon^{2} c^{2} / 2$ edge deletions occur in $\kappa(G)$ during the phase. To complete the proof, we will show that the corresponding number of edge insertions in $\kappa(G)$ is also $O\left(\epsilon^{2} c^{2}\right)$.

For an edge insertion in $G$, at most one edge is inserted into $\kappa(G)$. For an edge deletion in $G$, there can be $O(\epsilon c)$ edge insertions in $\kappa(G)$, but only if the subroutine REFILL(.) is called. Lemma 5.5 shows that at most $\epsilon c$ calls are made to the subroutine $\operatorname{REFILL}($.$) during the phase. This implies that at most O\left(\epsilon^{2} c^{2}\right)$ edge insertions occur in $\kappa(G)$ during the phase, which proves the second part of the theorem.

Lemma 5.5. The subroutine REFILL(.) is called at most $\epsilon c$ times during the phase.
Proof. When the phase begins, every tight node has exactly $c$ friends (see Theorem 3.15), and during the phase, the status of a node $u$ is changed from slack to tight only when $\left|\kappa\left(\mathcal{N}_{u}\right)\right|$ becomes at least $c$. On the other hand, the subroutine $\operatorname{REFILL}(u)$ is called only when $u$ is tight and $\left|\kappa\left(\mathcal{N}_{u}\right)\right|$ falls below $(1-\epsilon) c$. Thus, each call to REFILL(.) corresponds to a scenario where a tight node has lost at least $\epsilon c$ friends. Each edge deletion in $G$ leads to at most two such losses (one for each of the endpoints), whereas an edge insertion in $G$ leads to no such event. Since the phase lasts for $\epsilon^{2} c^{2} / 2$ edge insertions/deletions in $G$, a counting argument shows that it can lead to at most $\left(\epsilon^{2} c^{2} / 2\right) \cdot 2 /(\epsilon c)=\epsilon c$ calls to REFILL(.).

```
IF the list neighbors \((u)\) is empty, Then
    Set TyPE \((u) \leftarrow\) slack.
    Return.
Let \(x\) be the first node in the list neighbors \((u)\).
While (\#Friends \((u)<c\) )
    If \(x \notin\) Friends \((u)\), Then
        Add \(x\) to \(\operatorname{Friends}(u)\) and \(u\) to \(\operatorname{Friends}(x)\), so that the edge \((u, x)\)
        becomes part of the kernel \(\kappa(G)=(V, \kappa(E))\). Increment each of the counters
        \#Friends \((u)\), \#Friends \((v)\) by one unit. Update the pointers
        Pointer \([u, x]\) and Pointer \([x, u]\).
    If \(x\) is the last node in NEIGHBors \((u)\), Then
        exit the While loop.
    Else
        Let \(y\) be the node that succeeds \(x\) in the list neighbors \((u)\).
        Set \(x \leftarrow y\).
IF \#Friends \((u)<c\), Then
    Set TyPE \((u) \leftarrow\) slack.
```

Figure 16: REFILL(u).

It remains to prove the final part of the theorem, which states that the algorithm maintains an $(\epsilon, c)$ kernel. Specifically, we will show that throughout the duration of the phase, the subgraph $\kappa(G)$ maintained by the algorithm satisfies Invariants 3.6-3.8.

Lemma 5.6. Suppose that $\kappa(G)$ satisfies Invariants 3.7, 3.8 before an edge update in $G$. Then these invariants continue to hold even after we modify $\kappa(G)$ as per the procedure in Section 5.2.1 (resp. Section 5.2.2).

Proof. Follows from the descriptions of the procedures in Sections 5.2.1 and 5.2.2.
Recall that in the beginning of the phase, the graph $\kappa(G)$ is a $(0, c)$-kernel of $G$. Since every $(0, c)$ kernel is also an $(\epsilon, c)$-kernel, we repeatedly invoke Lemma 5.6 after each update in $G$, and conclude that $\kappa(G)$ satisfies Invariants 3.7, 3.8 throughout the duration of the phase. For the rest of this section, we focus on proving the remaining Invariant 3.6.

Fix any node $v \in V$. When the phase begins, the subgraph $\kappa(G)$ is a $(0, c)$-kernel of $G$, so that we have $\left|\kappa\left(\mathcal{N}_{v}\right)\right| \leq c$. During the phase, the node $v$ can get new friends under two possible situations.

1. An edge incident to $v$ has just been inserted into (resp. deleted from) the graph $G$, and the procedure in Section 5.2.1 (resp. the subroutine $\operatorname{REFILL}(v)$ in Figure 16) is going to be called.
2. The subroutine $\operatorname{REFILL}(u)$ is going to be called for some $u \in \mathcal{N}_{v}$.

If we are in situation (1) and the node $v$ already has at least $c$ friends, then the procedure under consideration will not end up adding any more node to $\kappa\left(\mathcal{N}_{v}\right)$. Thus, it suffices to show that the node $v$ can get at most $\epsilon c$ new friends during the phase under situation (2). Note that each call to $\operatorname{REFILL}(u), u \neq v$, creates at most one new friend for $v$. Accordingly, if we show that the subroutine $\operatorname{REFILL}($.$) is called at$ most $\epsilon c$ times during the entire phase, then this will suffice to conclude the proof of Theorem 5.2. But this has already been done in Lemma 5.5.

## $6(4+\epsilon)$-approximate matching in $O\left(m^{1 / 3} / \epsilon^{2}\right)$ worst-case update time

Fix any $\epsilon \in(0,1 / 6)$. We present an algorithm for maintaining a $(4+\epsilon)$-approximate matching $M$ in a dynamic graph $G=(V, E)$ with $O\left(m^{1 / 3} / \epsilon^{2}\right)$ worst-case update time.

### 6.1 Overview of our approach.

As in Section 5.1, we partition the sequence of updates (edge insertions/deletions) in $G$ into phases. Each phase lasts for $\epsilon^{2} c^{2} / 2$ consecutive updates in $G$. Let $G_{i, t}$ denote the state of $G$ just after the $t^{\text {th }}$ update in phase $i$, where $0 \leq t \leq \epsilon^{2} c^{2} / 2$. The initial state of the graph, before it starts changing, is given by $G_{1,0}$. Thus, we reach the graph $G_{i, t}$ from $G_{1,0}$ after a sequence of $(i-1) \cdot\left(\epsilon^{2} c^{2} / 2\right)+t$ updates in $G$.

For the rest of this section, we focus on describing our algorithm for any given phase $i \geq 1$. We define $m \leftarrow\left|E_{i, 0}\right|$ to be the number of edges in the input graph in the beginning of the phase, and set $c=m^{1 / 3}$. Since the phase lasts for only $O\left(\epsilon^{2} m^{2 / 3}\right)$ updates in $G$, it follows that $\left|E_{i, t}\right|=O(m)$ for all $0 \leq t \leq \epsilon^{2} c^{2} / 2$.
Preprocessing. If $i=1$, then we build a $(0, c)$-kernel $\kappa\left(G_{1,0}\right)$ in the beginning of the phase as per Theorem 3.15. Next, we build a maximal matching $M_{1,0}$ in $\kappa\left(G_{1,0}\right)$. Overall, this takes $O\left(|V|+\left|E_{1,0}\right|\right)$ time.

Algorithm for each phase. We first present the high level approach. The main difference between the framework of Section 5.1 and that of this section is as follows. In Section 5.1, the algorithm rebuilds the kernel and a corresponding matching from scratch at the beginning of each phase, and we amortize the cost of the rebuild over the phase. In this section, to achieve a worst-case running time we do this rebuilding "in the background" during the phase. This means that at the beginning of a phase we start with an empty graph $G^{*}$ and an empty kernel and insert $O\left(m /\left(\epsilon^{2} c^{2}\right)\right)$ edges into $G^{*}$ and its kernel during each update in $G$. Thus, edge insertions into $G^{*}$ need to be handled in constant time. We can handle "bunch updates" into $G^{*}$ to build the kernel efficiently, but we cannot efficiently update a $3 / 2$-approximate matching in the kernel as each edge insertion takes $O(c)$ time (Theorem 3.4). However, we can update a maximal matching, i.e., a 2 -approximate matching in constant time per edge insertions (Theorem 3.1). Thus, we run the 2approximation algorithm on the kernel instead of the $3 / 2$-approximation algorithm, leading to a $(4+\epsilon)$ overall approximation. Additionally, we need to perform the updates of the current phase in $G^{*}$. For that we use basically the same algorithm as in Section 5.2. As a result at the end of a phase $G^{*}=G$ and the kernel that we built for $G^{*}$ is only an $(\epsilon, c)$ kernel of $G$, i.e., we have to start the next phase with a $(\epsilon, c)$-kernel instead of a $(0, c)$-kernel. This, however, degrades the approximation ratio only by an additional factor of $\epsilon$.

To summarize our algorithm for phase $i$ has two components.

- Dealing with the current phase. Just before the start of phase $i$, an $(\epsilon, c)$-kernel $\kappa\left(G_{i, 0}\right)$ and a maximal matching $M_{i, 0}$ in $\kappa\left(G_{i, 0}\right)$ are handed over to us. Then, as $G$ keeps changing, we keep modifying the subgraph $\kappa(G)$ and the matching $M$. Till the end of phase $i$, we ensure that $\kappa(G)$ remains a $(2 \epsilon, c)$ kernel of $G$ and that $M$ remains a maximal matching in $\kappa(G)$. Hence, by Theorem 3.9, the matching $M$ gives a $(4+\epsilon)$-approximation ${ }^{6}$ to the maximum matching in $G$ throughout the duration of phase $i$.
- Preparing for the next phase. We build a new $(\epsilon, c)$-kernel (and a maximal matching in it) for $G_{i+1,0}$ in the background. They are handed over to the algorithm for phase $(i+1)$ at the start of phase $(i+1)$.
We elaborate upon each of these components in more details in Section 6.2 (see Lemmas 6.4, 6.5). For maintaining a kernel in a dynamic graph, both these components use a procedure that is essentially the one described in Section 5.2. However, a more fine tuned analysis of the procedure becomes necessary.

Finally, Theorem 3.9 and Lemmas 6.4, 6.5 give us the desired guarantee.

[^5]Theorem 6.1. We can maintain a $(4+\epsilon)$-approximate matching in a dynamic graph $G=(V, E)$ in $O\left(m^{1 / 3} / \epsilon^{2}\right)$ worst-case update time.

### 6.2 Algorithm for maintaining and building the kernel during each phase.

In Subsections 6.2.3 and 6.2.4 we present an algorithm for maintaining a kernel in a fully dynamic graph. It is basically a relaxation of the algorithm from Subsection 5.2 and is based on the concept of an "epoch".

Definition 6.2. For $c \geq 1, \lambda \in(0,1)$, a sequence of updates in a graph is called a $(\lambda, c)$-epoch iff it contains at most $\lambda^{2} c^{2} / 2$ edge deletions (which can be arbitrarily interspersed with any number of edge insertions).

The resulting algorithm is summarized in the theorem below.
Theorem 6.3. Fix any $c \geq 1$ and $\epsilon, \lambda>0$ with $(\epsilon+\lambda)<1 / 3$. Consider any $(\lambda, c)$-epoch in the input graph $G=(V, E)$, and suppose that we are given an $(\epsilon, c)$-kernel $\kappa(G)$ in the beginning of the epoch. Then we have an algorithm for updating $\kappa(G)$ after each update in $G$. The algorithm satisfies three properties.

1. An edge insertion in $G$ is handled in $O(1)$ worst-case time, and this leads to at most one edge insertion in $\kappa(G)$ and zero edge deletion in $\kappa(G)$.
2. An edge deletion in $G$ is handled in $O(c)$ worst-case time, and this leads to at most $O(\lambda c)$ edge insertions in $\kappa(G)$ and at most one edge deletion in $\kappa(G)$.
3. Throughout the duration of the epoch, $\kappa(G)$ remains $a(\lambda+\epsilon, c)$-kernel of $G$.

We use the kernel update algorithm in the two components of a phase, which we describe next in detail. Recall that $G_{i, t}$ denotes the state of $G$ just after the $t^{t h}$ update in phase $i \geq 1$ for all $0 \leq t \leq \epsilon^{2} c^{2} / 2$, $m=\left|E_{i, 0}\right|$, and $c=m^{1 / 3}$.

### 6.2.1 Dealing with the current phase

In the beginning of phase $i$, we are given an $(\epsilon, c)$-kernel $\kappa\left(G_{i, 0}\right)$ and a maximal matching $M_{i, 0}$ in $\kappa\left(G_{i, 0}\right)$. Since a phase lasts for $\epsilon^{2} c^{2} / 2$ edge updates in $G$, we treat phase $i$ as an $(\epsilon, c)$-epoch (see Definition 6.2). Thus, after each edge insertion/deletion in $G$ during phase $i$, we modify $\kappa(G)$ as per Theorem 6.3. This ensures that $\kappa(G)$ remains an $(2 \epsilon, c)$-kernel till the end of phase $i$. After each update in $\kappa(G)$, we modify the matching $M$ using Theorem 3.1 so as to ensure that $M$ remains a maximal matching in $\kappa(G)$.

We show that this procedure requires $O(c)$ time in the worst case to handle an update in $G$.

- Case 1: An edge is inserted into the graph $G$. By Theorem 6.3, in this case updating the kernel $\kappa(G)$ requires $O(c)$ time, and this results in at most one edge insertion into $\kappa(G)$ and zero edge deletion from $\kappa(G)$. By Theorem 3.1, updating the matching $M$ requires $O(1)$ time.
- Case 2: An edge is deleted from the graph $G$. By Theorem 6.3, in this case updating the kernel $\kappa(G)$ requires $O(c)$ time, and this results in $O(\epsilon c)$ edge insertions into $\kappa(G)$ and at most one edge deletion from $\kappa(G)$. Since the maximum degree of a node in $\kappa(G)$ is $O(c)$, updating the matching $M$ requires $O(\epsilon c)+O(c)=O(c)$ time (see Theorem 3.1).

Lemma 6.4. Suppose that when phase $i$ begins, we are given an $(\epsilon, c)$-kernel $\kappa(G)$ of $G$ and a maximal matching $M$ in $\kappa(G)$. After each update in $G$ during phase $i$, we can modify $\kappa(G)$ and $M$ in $O(c)$ time. Till the end of phase $i, \kappa(G)$ remains an $(2 \epsilon, c)$-kernel of $G$ and $M$ remains a maximal matching in $\kappa(G)$.

### 6.2.2 Preparing for the next phase.

Before phase $i$ begins, in $O(1)$ time we initialize a graph $G^{*}=\left(V, E^{*}\right)$ with $E^{*}=\emptyset$, a $(0, c)$-kernel $\kappa\left(G^{*}\right)=\left(V, \kappa\left(E^{*}\right)\right)$, and a maximal matching $M^{*}=\emptyset$ in $\kappa\left(G^{*}\right)$. To achieve constant time we use an uninitialized array for the TyPE-bit and assume that every node for which the bit is not initialized is of type slack.

```
IF an edge \((u, v)\) is inserted into \(G\), THEN
    \(E^{*} \leftarrow E^{*} \cup\{(u, v)\}\).
ELSE IF an edge \((u, v)\) is deleted from \(G\), THEN
    \(E^{*} \leftarrow E^{*} \backslash\{(u, v)\}\).
FOR \(i=1\) to \(2 m /\left(\epsilon^{2} c^{2}\right)\)
    \(\operatorname{IF} E \backslash E^{*} \neq \emptyset\), Then
    Pick any edge \(e \in E \backslash E^{*}\).
    \(E^{*} \leftarrow E^{*} \cup\{e\}\).
```

Figure 17: Update- $G^{*}$.

After each update in $G$ in phase $i$, we call Update- $G^{*}$ (Figure 17). This ensures the following properties.

- (P1) Each edge insertion in $G$ leads to at most $O\left(m /\left(\epsilon^{2} c^{2}\right)\right)$ edge insertions in $G^{*}$.
- (P2) Each edge deletion in $G$ leads to at most one edge deletion in $G^{*}$ and at most $O\left(m /\left(\epsilon^{2} c^{2}\right)\right)$ edge insertions in $G^{*}$.
- (P3) After each edge insertion/deletion in $G$, changing the graph $G^{*}$ requires $O\left(m /\left(\epsilon^{2} c^{2}\right)\right)$ time.
- (P4) We always have $E^{*} \subseteq E$. Furthermore, at the end of phase $i$, we have $E^{*}=E$. This holds since $m$ is the size of $E$ in the beginning of phase $i$, and since the phase lasts for $\epsilon^{2} c^{2} / 2$ updates in $G$.

Now, consider the sequence of updates in $G^{*}$ that take place during phase $i$. (P1) and (P2) guarantee that this sequence contains at most $\epsilon^{2} c^{2} / 2$ edge deletions. So this sequence can be treated as an $(\epsilon, c)$-epoch in $G^{*}$ (see Definition 6.2). Thus, after each update in $G^{*}$, we modify the subgraph $\kappa\left(G^{*}\right)$ using Theorem 6.3, and subsequently, we modify the maximal matching $M^{*}$ in $\kappa\left(G^{*}\right)$ using Theorem 3.1.

Since $\kappa\left(G^{*}\right)$ was a $(0, c)$-kernel of $G^{*}$ in the beginning of the epoch, we get the following guarantee.

- (P5) Throughout the duration of phase $i$, the graph $\kappa\left(G^{*}\right)$ is an $(\epsilon, c)$-kernel of $G^{*}=\left(V, E^{*}\right)$ and $M^{*}$ is a maximal matching in $\kappa\left(G^{*}\right)$.

By (P1) and (P2), an update in $G$ leads to $O\left(m /\left(\epsilon^{2} c^{2}\right)\right)$ edge insertions and at most one edge deletion in $G^{*}$. Each of these edge insertions in $G^{*}$ further leads to at most one edge insertion in $\kappa\left(G^{*}\right)$, while the potential edge deletion in $G^{*}$ leads to at most one edge deletion and $O(\epsilon c)$ edge insertions in $\kappa\left(G^{*}\right)$ (see Theorem 6.3). To summarize, an update in $G$ leads to the following updates in $G^{*}$ and $\kappa\left(G^{*}\right)$.

- At most one edge deletion and $O\left(m /\left(\epsilon^{2} c^{2}\right)\right)$ edge insertions in $G^{*}$.
- At most one edge deletion and $O\left(\epsilon c+m /\left(\epsilon^{2} c^{2}\right)\right)$ edge insertions in $\kappa\left(G^{*}\right)$.

Thus, by Theorems 6.3 and 3.1, we get the following bound on the update time

- (P6) After each edge update in $G$, the graph $\kappa\left(G^{*}\right)$ and the maximal matching $M^{*}$ in $\kappa\left(G^{*}\right)$ can be modified in $O(c)+O\left(\epsilon c+m /\left(\epsilon^{2} c^{2}\right)\right)=O\left(c+m /\left(\epsilon^{2} c^{2}\right)\right)$ time.

Using (P3) and (P6), we derive that after each update in $G$, the time required to modify the structures $G^{*}, \kappa\left(G^{*}\right)$ and $M^{*}$ is $O\left(c+m /\left(\epsilon^{2} c^{2}\right)\right)=O\left(m^{1 / 3} / \epsilon^{2}\right)$ in the worst case. By (P4) and (P5), the graph $G^{*}$ becomes identical with $G$ at the end of phase $i$, and at this point $\kappa\left(G^{*}\right)$ becomes an $(\epsilon, c)$-kernel of $G_{i+1,0}$.

Lemma 6.5. Starting from the beginning of phase $i$, we can run an algorithm with the following properties.

- After each update in $G$ during phase i, it performs $O\left(\mathrm{~m}^{1 / 3} / \epsilon^{2}\right)$ units of computation.
- At the end of phase $i$, it returns an $(\epsilon, c)$-kernel of $G_{i+1,0}$ and a maximal matching in this kernel.

Theorem 6.1 now follows from Theorem 3.9 and Lemmas 6.4, 6.5.

### 6.2.3 Handling an edge insertion into $G$ during the epoch.

Suppose that an edge $(u, v)$ is inserted into the graph $G=(V, E)$. To handle this edge insertion, we first update the lists NEighbors $(u)$ and NEighbors $(v)$, and then process the edge as follows.

- Case 1: $\operatorname{Either} \operatorname{Type}(u)=$ tight or $\operatorname{Type}(v)=$ tight.

We do nothing and conclude the procedure.

- Case 2: Both $\operatorname{Type}(u)=$ slack and $\operatorname{TyPE}(v)=$ slack.
- Case 2a: Either \#Friends $(u) \geq(1-\epsilon) c$ or \#Friends $(v) \geq(1-\epsilon) c$. If \#Friends $(u) \geq(1-\epsilon) c$, then we set $\operatorname{Type}(u) \leftarrow$ tight. Next, if \#Friends $(v) \geq(1-\epsilon) c$, then we set TyPE $(v) \leftarrow$ tight.
- Case 2b: Both \#Friends $(u)<(1-\epsilon) c$ and \#Friends $(v)<(1-\epsilon) c$.

We add the edge $(u, v)$ to the kernel $\kappa(G)$. Specifically, we add $u$ to the list Friends $(v)$ and $v$ to the list Friends $(u)$, update the pointers Pointer $[u, v]$ and Pointer $[v, u]$ accordingly, and increment each of the counters \#Friends $(u)$, \#Friends $(v)$ by one unit.

Lemma 6.6. Suppose that an edge is inserted into the graph $G$ and we run the procedure described above.

1. This causes at most one edge insertion into $\kappa(G)$ and at most one edge deletion from $\kappa(G)$.
2. The procedure runs in $O(1)$ time in the worst-case.

### 6.2.4 Handling an edge deletion in $G$ during the epoch.

Suppose that an edge $(u, v)$ gets deleted from the graph $G=(V, E)$. To handle this edge deletion, we first consider the adjacency-list data structure for $G$, and update the lists NEIGHBORS $(u)$ and NEIGHBORS $(v)$.

We check if the edge $(u, v)$ was part of the kernel $\kappa(G)$, and, if yes, then we delete $(u, v)$ from $\kappa(G)$. Specifically, we delete $u$ from $\operatorname{Friends}(v)$ and $v$ from Friends ( $u$ ) (using the pointers Pointer $[u, v]$, Pointer $[v, u]$ ), and decrement each of the counters \#Friends $(u)$, \#Friends $(v)$ by one unit.

We then process the nodes $u$ and $v$ one after another. Below, we describe only the procedure that runs on the node $u$. The procedure for the node $v$ is exactly the same.

- Case 1: $\operatorname{Type}(u)=$ tight. Here, we check if the number of friends of $u$ has dropped below the prescribed limit due to the edge deletion, and, accordingly, we consider two possible sub-cases.
- Case 1a: \#Friends $(u)<(1-\lambda-\epsilon) c$. Here, we call the subroutine REFILL-NOW $(u)$ (see Figure 18).
- Case 1b: \#Friends $(u) \geq(1-\lambda-\epsilon) c$. In this case, we do nothing and conclude the procedure.
- Case 2: $\operatorname{Type}(u)=$ slack. In this case, we do nothing and conclude the procedure.

Remark. The procedure REFILL-NOW $(u)$ (see Figure 18) is identical with the procedure REFILL $(u)$ (see Figure 16) in Section 5, except that it uses a different threshold for the number of friends of a node.

Lemma 6.7. Suppose that an edge is deleted from the graph $G$ and we run the procedure described above.

1. This causes at most $O(\lambda c)$ edge insertions into $\kappa(G)$ and at most one edge deletion from $\kappa(G)$.
2. The procedure runs in $O(c)$ time in the worst-case.
```
IF the list Neighbors(u) is empty, Then
    Set Type(u)\leftarrow slack.
    REtURN.
Let }x\mathrm{ be the first node in the list NEighbors(u).
While (#Friends (u)< (1-\epsilon)c)
    IF }x\not\in\operatorname{FriendS}(u)\mathrm{ , Then
        Add x to Friends(u) and u}\mathrm{ to Friends(x), so that the edge (u,x) becomes part of the
        kernel }\kappa(G)=(V,\kappa(E)). Increment each of the counters #FriEnds(u), #Friends(x
        by one unit. Update the pointers Pointer [u,x], Pointer [x,u].
    IF}x\mathrm{ is the last node in NEIGHbors(u), Then
            exit the While loop.
    ElSE
            Let }y\mathrm{ be the node that succeeds }x\mathrm{ in the list NEighbors(u).
            Set }x\leftarrowy
IF #Friends (u)< (1-\epsilon)c, Then
    Set TyPE (u)\leftarrow slack.
```

Figure 18: REFILL-NOW $(u)$.

### 6.2.5 Proof of Theorem 6.3

The first and the second parts of the theorem follows from Lemma 6.6 and Lemma 6.7 respectively. It remains to show that the algorithm maintains a $(\lambda+\epsilon, c)$-kernel. Specifically, we will show that throughout the duration of the $(\lambda, c)$-epoch, the subgraph $\kappa(G)$ maintained by the algorithm continues to satisfy Invariants 3.6-3.8 (replacing $\epsilon$ by $\lambda+\epsilon$ in Invariants 3.6, 3.7).

Lemma 6.8. Suppose that $\kappa(G)$ satisfies two conditions before an edge insertion (resp. deletion) in $G$.

1. $\left|\kappa\left(\mathcal{N}_{v}\right)\right| \geq(1-\lambda-\epsilon)$ c for all nodes $v \in \kappa_{T}(V)$.
2. For all nodes $u, v \in \kappa_{S}(V)$, if $(u, v) \in E$, then either $u \in \kappa\left(\mathcal{N}_{v}\right)$ or $v \in \kappa\left(\mathcal{N}_{u}\right)$.

Then these two conditions continue to hold even after we modify $\kappa(G)$ using the procedure in Section 6.2.3 (resp. Section 6.2.4).

Proof. Follows from the descriptions of the procedures in Sections 6.2.3 and 6.2.4.
Recall that before the $(\lambda, c)$-epoch begins, the graph $\kappa(G)$ is an $(\epsilon, c)$-kernel of $G$. Since every $(\epsilon, c)$ kernel is also a $(\lambda+\epsilon, c)$-kernel, we repeatedly invoke Lemma 6.8 after each update in $G$, and conclude that $\kappa(G)$ satisfies Invariants 3.7, 3.8 (replacing $\epsilon$ by $\lambda+\epsilon$ ) throughout the duration of the epoch. For the rest of this section, we focus on proving the remaining Invariant 3.6. Specifically, we will show that at any point in time during the epoch, we have $\left|\kappa\left(\mathcal{N}_{v}\right)\right| \leq(1+\lambda+\epsilon) c$ for all $v \in V$.

Fix any node $v \in V$. When the $(\lambda, c)$-epoch begins, the subgraph $\kappa(G)$ is an $(\epsilon, c)$-kernel of $G$, so that we have $\left|\kappa\left(\mathcal{N}_{v}\right)\right| \leq(1+\epsilon) c$. During the epoch, the node $v$ can get new friends under two possible situations.

1. An edge incident to $v$ has just been inserted into (resp. deleted from) the graph $G$, and the procedure in Section 6.2 .3 (resp. the subroutine REFILL-NOW $(v)$ ) is going to be called.
2. The subroutine REFILL-NOW $(u)$ is going to be called for some $u \in \mathcal{N}_{v}$.

If we are in situation (1) and the node $v$ already has more than $(1+\epsilon) c$ friends, then the procedure under consideration will not end up adding any more node to $\kappa\left(\mathcal{N}_{v}\right)$. Thus, it suffices to show that the node $v$ can get at most $\lambda c$ new friends during the epoch under situation (2). Note that each call to REFILL-NOW $(u)$, $u \neq v$, creates at most one new friend for $v$. Accordingly, if we show that the subroutine REFILL-NOW(.) is called at most $\lambda c$ times during the entire epoch, then this will suffice to conclude the proof of Theorem 6.3.

Lemma 6.9. The subroutine REFILL-NOW(.) is called at most $\lambda c$ times during $a(\lambda, c)$-epoch.
Proof. When the epoch begins, every tight node has at least $(1-\epsilon) c$ friends, and during the epoch, the status of a node $u$ is changed from slack to tight only when $\left|\kappa\left(\mathcal{N}_{u}\right)\right|$ exceeds $(1-\epsilon) c$. On the other hand, the subroutine REFILL-NOW $(u)$ is called only when $u$ is tight and $\left|\kappa\left(\mathcal{N}_{u}\right)\right|$ falls below $(1-\lambda-\epsilon) c$. Thus, each call to REFILL-NOW(.) corresponds to a scenario where a tight node has lost at least $\lambda c$ friends. Each edge deletion in $G$ leads to at most two such losses (one for each of the endpoints), whereas an edge insertion in $G$ leads to no such event. Since a $(\lambda, c)$-epoch contains at most $\lambda^{2} c^{2} / 2$ edge deletions (Definition 6.2), a counting argument shows that such an epoch can result in at most $\left(\lambda^{2} c^{2} / 2\right) \cdot 2 /(\lambda c)=\lambda c$ calls to REFILL-NOW(.).

## References

[1] Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In 55th IEEE Symposium on Foundations of Computer Science, 2014. 1
[2] Surender Baswana, Manoj Gupta, and Sandeep Sen. Fully dynamic maximal matching in $O(\log n)$ update time. In 52nd IEEE Symposium on Foundations of Computer Science, pages 383-392, 2011. 1, 2, 3
[3] Ran Duan and Seth Pettie. Approximating maximum weight matching in near-linear time. In 51st IEEE Symposium on Foundations of Computer Science, pages 673-682, 2010. 1
[4] Manoj Gupta and Richard Peng. Fully dynamic ( $1+\epsilon$ )-approximate matchings. In 54th IEEE Symposium on Foundations of Computer Science, pages 548-557, 2013. 1, 2, 3
[5] John E. Hopcroft and Richard M. Karp. An $n^{5 / 2}$ algorithm for maximum matchings in bipartite graphs. SIAM Journal on Computing, 2(4):225-231, 1973. 1, 20
[6] Zoran Ivkovic and Errol L. Lloyd. Fully dynamic maintenance of vertex cover. In Graph-Theoretic Concepts in Computer Science, pages 99-111, 1993. 1
[7] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within $2-\epsilon$. Journal of Computer and System Sciences, 74(3):335-349, 2008. 1
[8] László Lovász and Michael D. Plummer. Matching Theory. Akadémiai Kiadó, Budapest, 1986. Also published as Vol. 121 of the North-Holland Mathematics Studies, North Holland Publishing, Amsterdam. 2
[9] Aleksander Madry. Navigating central path with electrical flows: From flows to matchings, and back. In 54th IEEE Symposium on Foundations of Computer Science, pages 253-262, 2013. 1
[10] Silvio Micali and Vijay V. Vazirani. An $O(\sqrt{|V|}|E|)$ algorithm for finding maximum matching in general graphs. In 21st IEEE Symposium on Foundations of Computer Science, pages 17-27, 1980. 1
[11] Marcin Mucha and Piotr Sankowski. Maximum matchings via Gaussian elimination. In 45th IEEE Symposium on Foundations of Computer Science, pages 248-255, 2004. 1
[12] Ofer Neiman and Shay Solomon. Simple deterministic algorithms for fully dynamic maximal matching. In 45th ACM Symposium on Theory of Computing, pages 745-754, 2013. 1, 2, 3,20
[13] Krzysztof Onak and Ronitt Rubinfeld. Maintaining a large matching and a small vertex cover. In 42nd ACM Symposium on Theory of Computing, pages 457-464, 2010. i, 1, 2, 3, 6, 7
[14] Piotr Sankowski. Faster dynamic matchings and vertex connectivity. In 18th ACM-SIAM Symposium on Discrete Algorithms, pages 118-126, 2007. 1


[^0]:    *An extended abstract of this paper, not containing the algorithm in Section 4 and not containing the proof of Theorem 2.11, has been accepted for publication in ACM-SIAM Symposium on Discrete Algorithms (SODA)' 2015.
    ${ }^{\dagger}$ Email: bsayan@imsc.res.in. The Institute of Mathematical Sciences, Chennai, India. This work was done while the author was in Faculty of Computer Science, University of Vienna, Austria. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programe (FP7/2007-2013) / ERC Grant Agreement number 340506.
    ${ }^{\ddagger}$ Email: monika.henzinger@univie.ac.at. Faculty of Computer Science, University of Vienna, Austria. The research leading to these results has received funding from the European Unions Seventh Framework Programme (FP7/2007-2013) under grant agreement 317532 and from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Grant Agreement number 340506.
    ${ }^{\text {§ }}$ Email: giuseppe.italiano@uniroma2.it. Università di Roma "Tor Vergata", Rome, Italy. Partially supported by MIUR, the Italian Ministry of Education, University and Research, under Project AMANDA (Algorithmics for MAssive and Networked DAta).

[^1]:    ${ }^{1}$ The proof actually shows a stronger result assuming that the level change of node $v$ from $i$ to $i-1$ causes $\Theta\left(D_{v}(0, i)\right)$ many edges to change their level.

[^2]:    ${ }^{2}$ We write $\Theta\left(1+D_{v}(0, k)\right)$ instead of $\Theta\left(D_{v}(0, k)\right)$ as $D_{v}(0, k)$ can be zero.

[^3]:    ${ }^{3}$ A passive node is not adjacent to any edges and, thus, its weight is zero and it has no neighbors.
    ${ }^{4}$ Note that $D_{v}(0, k)$ may be zero. Hence, we bound the runtime of MOVE-UP $(v)$ by $\Theta\left(1+D_{v}(0, k)\right)$ instead of $\Theta\left(D_{v}(0, k)\right)$.

[^4]:    ${ }^{5}$ Since $D_{v}(0, k)$ may be zero, we bound the runtime of MOVE-DOWN $(v)$ by $\Theta\left(1+D_{v}(0, k)\right)$ instead of $\Theta\left(D_{v}(0, k)\right)$.

[^5]:    ${ }^{6}$ To be precise, the theorem shows a $(4+12 \epsilon)$-approximation but running the algorithm with $\epsilon^{\prime}=\epsilon / 12$ results in a $(4+\epsilon)$ approximation.

