

Online bipartite matching in offline time

Bartłomiej Bosek*, Dariusz Leniowski†, Piotr Sankowski† and Anna Zych†

*Theoretical Computer Science Department, Faculty of Mathematics and Computer Science,
Jagiellonian University, Kraków, Poland, Email: bosek@tcs.uj.edu.pl

†Institute of Computer Science, University of Warsaw, Warszawa, Poland, Email: {d.leniowski,sank,anka}@mimuw.edu.pl

Abstract—This paper investigates the problem of maintaining maximum size matchings in incremental bipartite graphs. In this problem a bipartite graph G between n clients and n servers is revealed online. The clients arrive in an arbitrary order and request to be matched to a subset of servers. In our model we allow the clients to switch between servers and want to maximize the matching size between them, i.e., after a client arrives we find an augmenting path from a client to a free server. Our goals in this model are twofold. First, we want to minimize the number of times clients are reallocated between the servers. Second, we want to give fast algorithms that recompute such reallocation.

As for the number of changes, we propose a greedy algorithm that chooses an augmenting path π that minimizes the maximum number of times each server in π was used by augmenting paths so far. We show that in this algorithm each server has its client reassigned $O(\sqrt{n})$ times. This gives an $O(n^{3/2})$ bound on the total number of changes, what gives a progress towards the main open question risen by Chaudhuri *et al.* (INFOCOM'09) who asked to prove $O(n \log n)$ upper bound. Next, we argue that the same bound holds in the decremental case. Moreover, we show incremental and decremental algorithms that maintain $(1 - \varepsilon)$ -approximate matching with total of $O(\varepsilon^{-1}n)$ reallocations, for any $\varepsilon > 0$.

Finally, we address the question of how to efficiently compute paths given by this greedy algorithm. We show that by introducing proper amortization we can obtain an incremental algorithm that maintains the maximum size matching in total $O(\sqrt{nm})$ time. This matches the running time of one of the fastest static maximum matching algorithms that was given by Hopcroft and Karp (SIAM J. Comput '73). We extend our result to decremental case where we give the same total bound on the running time. Additionally, we show $O(\varepsilon^{-1}m)$ time incremental and decremental algorithms that maintain $(1 - \varepsilon)$ -approximate matching for any $\varepsilon > 0$. Observe that this bound matches the running time of the fastest approximate static solution as well.

Index Terms—online matchings; bipartite matchings; incremental and decremental algorithms; approximate matchings;

I. INTRODUCTION

In this paper, we study an incremental bipartite matching problem, which corresponds to a scenario in which clients arrive online and request service from a set of given servers. Each client arrives together with a set of servers being able to handle his request. In this scenario we want to provide the service to as many clients as possible. Hence, when necessary we reallocate the clients between the servers, and the cost of the solution is the total number of times a client is allocated to a server during the course of this reallocation process. The solution to this problem needs to find an augmenting

path (a path that alternates between matched and unmatched edges) in the current graph from the client to some unused server. Using this language we can define, in an equivalent way, the cost to be equal to the total length of all augmenting paths. This approach contrasts with the model introduced by Karp, Vazirani and Vazirani in [14], where clients are assigned to servers permanently. Our model was introduced by Grove *et al.* [8] in '95, where the authors applied competitive analysis to show an $O(\log n)$ bound when each client was connecting to at most two servers. Next, the model was studied by Chaudhuri *et al.* [7] where a total bound of $O(n \log n)$ was shown in some very restricted models, e.g., forests or random graphs with degree $\Theta(\log n)$. These special cases led the authors to conjecture that a greedy algorithm that chooses the shortest augmenting path achieves a bound of $O(n \log n)$. However, they were unable to make any progress in the general case where only the trivial bound of $O(n^2)$ is known. Hence, our result constitutes a step towards resolving this main open problem from [7]. Here, we propose a new greedy algorithm for this problem. It uses the notion of a rank of a server-vertex, which denotes the number of times the server was used by an augmenting path so far. In each step it chooses an augmenting path π that minimizes the maximum rank of vertices on π . We prove that this algorithm reallocates each client at most $O(n^{1/2})$ times which gives the total bound of $O(n^{3/2})$.

As argued in [7] the above model has a wide set of applications in very diverse areas of computer science, e.g., streaming content delivery, web hosting, remote data storage, job scheduling, hashing, etc. In all these applications we have some clients (jobs for scheduling, or objects for hashing) arriving online that have to be allocated to some fixed set of servers (table cells for hashing). In all these cases we do not want to drop clients from service, but prefer to reallocate them so that as many clients as possible are satisfied. Here, the cost of not serving a client is much higher than the cost of reallocating the client to another server (e.g., copying an object in a hash table is cheap). Still we would like to minimize the cost of these changes. It is somewhat astonishing that despite the usefulness of this model no better solution than the naive $O(n^2)$ one was known till now.

Besides solving the incremental case, we extend our results in a twofold way. First, we show that the same bound holds in the decremental case, what follows by an easy reduction. Second, we show an incremental as well as decremental algorithm that is able to maintain $(1 - \varepsilon)$ -approximate matching using $O(\varepsilon^{-1}n)$ changes, for any $\varepsilon > 0$. This shows that by

This work was partially supported by ERC StG project PAAI 259515, FET IP project MULTIPLEX 317532, NCN grant 2013/11/D/ST6/03100 and N206 567940.

dropping an arbitrary small, but constant fraction of clients, an optimal linear bound on the number of changes can be obtained. Previously, in this incremental approximate model the best solution was either to use the algorithm of [14] to get $(1 - 1/e)$ -approximation in $O(n)$ reallocations, or to maintain exact matching in $O(n^2)$ reallocations.

So far we have been only counting the number of changes to the matching. Let us now change the perspective a bit and let us now consider the time needed to find these augmenting paths. Observe that our greedy algorithm can be implemented using bottleneck shortest paths, which can be found in $O(m + n \log n)$ time using Dijkstra's algorithm. Hence, a naive implementation of our algorithm would be very inefficient as it would require $O(nm + n^2 \log n)$.¹ However, our algorithm has an important property that the greedy paths it finds can be defined locally. In other words, we can show that no global information (e.g., shortest path distances) needs to be maintained in the graph to find the greedy paths. This allows us to devise a path search procedure that amortizes the required work in an efficient way. Intuitively the idea is to assign to each fixed vertex a number called *rank* that will only increase. Each time we try to look for a path that minimizes the ranks, and if we do not find one we raise them by one. We know that there always exists an augmenting path of rank bounded by $O(\sqrt{n})$, so either we find a path or exceed this bound and know that there is no augmenting path. This leads to incremental and decremental algorithms that maintain maximum matching in total $O(\sqrt{nm})$ time. Previously, only the trivial $O(nm)$ time solution was known. With this respect it is somewhat surprising that our algorithm finds augmenting paths one by one without looking on the whole graph when doing so. In particular, our running time matches one of the best offline solutions to the problem [10], but within the same bound computes not only one maximum matching, but n of them. Our algorithm can be extended to maintain $(1 - \varepsilon)$ -approximate matching in $O(\varepsilon^{-1}m)$ total time, for any $\varepsilon > 0$. Again this solution improves over the $O(m)$ time $(1 - 1/e)$ -approximate algorithm implied by [14] and matches the running time of the best offline solution.²

Finally we extend the algorithm to graphs with integer edge weights bounded by W , where it is able to maintain maximum weight matchings in total $O(W^{3/2}\sqrt{nm})$ time, or $(1 - \varepsilon)$ -approximation of them in total $O(W\varepsilon^{-1}m)$ time.

The bipartite matching problem is a central problem in combinatoric optimization and has a wide variety of applications, as it models the paradigmatic case of allocating objects to agents. Hence, having a tool to maintain such allocations in an incremental case is a very useful instrument. With this respect we note that our algorithms are purely combinatorial and very easy to implement. Implementation wise, our solution in some aspects is easier than the Hopcroft-Karp algorithm. We use only one graph search procedure (e.g., DFS or BFS),

¹The same bound holds for the algorithm proposed in [7], but there seem to be no possibility to improve it.

²It is folklore that running $O(\varepsilon^{-1})$ -phases of the Hopcroft-Karp algorithm results in $(1 - \varepsilon)$ -approximate matching and takes $O(\varepsilon^{-1}m)$ time.

whereas Hopcroft-Karp algorithm uses alternative calls to both DFS and BFS. Efficiency wise, our algorithms should perform in practice in a similar way as Hopcroft-Karp algorithm. We are not going to review all possible applications here, let us just mention one that gained a lot of attention recently, i.e., the AdWords problem [15]. In this application we indeed need to take care of vertices that arrive online. Moreover, we note that the algorithms that allocate AdWords might be guided by a solution constructed incrementally from previous requests. However, due to the large data size it is impossible to recompute each time the optimal solution, and only $(1 - 1/e)$ -approximate allocation can be used [16]. Hence, our framework could deliver tools that would allow to maintain better, or even optimal, solution in this application.

a) Related Work: Our work is most closely related to two lines of research. The first one, as discussed above, is the study of online matchings with augmentations [8], [7]. The second one is the recent series of papers on dynamic matchings [12], [20], [18], [4], [17], [9], [2]. One should note, however, that in all these papers a different dynamic model is studied. First, these papers consider edge updates (insertion and deletion of edges), whereas we consider vertex updates. As pointed above this model is closer to some applications than edge updates scenario, this was pointed as well in [5], [6], where it was noted that vertex-update dynamic problems are often more challenging. Second, all of them consider a more general fully-dynamic model, where the sequence of updates consists of intermixed insert and delete operations. When comparing these results with our incremental or decremental ones we derive the total time needed to handle m edge updates using these approaches. We first note that the solutions given in [12], [20] lose against the $O(nm)$ time algorithms, that is obtained by executing a search for an augmenting path each time a vertex is added or removed. On the other hand, the solution of [18] is superseded by [4] both in the running time and approximation factor. The solution of [4] results in a decremental as well incremental $1/2$ -approximate algorithm working in total of $O(m \log n)$ time. Here, we improve this by giving $(1 - \varepsilon)$ -approximation in $O(\varepsilon^{-1}m)$ total time. As for both [17], [9], the algorithm in [9] gives a better approximation than [17] in essentially the same running time. When translated into our model these result imply an algorithm that maintains $(1 - \varepsilon)$ -approximate matching in $O(\varepsilon^{-2}m\sqrt{m})$ total time. Our results compare well by giving either the same guarantee in much smaller $O(\varepsilon^{-1}m)$ total time, or by allowing to maintain exact maximum in better $O(\sqrt{nm})$ time. As for the weighted case none of the papers [17], [9], [2], where the problem is studied, is able to dynamically maintain exact maximum weight matching even when weights are bounded by a constant. On the other hand, we note that [9] implies an $O(W\varepsilon^{-2}\sqrt{mm})$ time algorithm that maintains $(1 - \varepsilon)$ -approximate matching. Here, again our $O(W\varepsilon^{-1}m)$ time algorithm gives a visible improvement.

Recently there was some progress on showing lower bound on the running time of dynamic algorithms [19], [1]. The second of these papers basing on some conjectures shows

polynomial lower bounds on the running time of incremental and decremental maximum matching algorithms. However, these results do not apply to our model, as they work with edge updates. There is no known reduction going in the required direction, because intuitively it would require to add vertices on both sides of the graph.

Finally, we note that our model is related to the online load balancing of permanent task with preemption. This model has been studied in many papers (see [3] for a survey). However, in our model we assume hard capacity constraints, whereas in the load balancing models one usually assumes that the capacities are soft, i.e., the capacities can be exceeded and one is interested in minimizing the maximum load. Hence, these results have somewhat different objectives.

b) Paper Overview: In the next section we formally define our problem, and introduce the basic framework that will be employed by our greedy approach. In Section III we prove some basic properties of ranks in this framework, whereas in Section IV basing on these properties we give the $O(\sqrt{n})$ upper bound on maximum ranks. Next, in Section V we give all the details of the amortized implementation of our algorithms. In Section VI we describe extensions of our result to approximate, as well as, to decremental and weighted scenario. Finally, Section VII concludes the paper and gives some interesting open problems.

II. PRELIMINARIES

In this section we formally define the online maximum matching problem in bipartite graphs and introduce the general framework we use to approach this problem. Let W and B be two sets of vertices over which the bipartite graph will be formed. The set W is given up front to the algorithm, i.e., corresponds to servers, whereas the vertices in B (clients) arrive online. Throughout the paper, the vertices of W are referred to as white and the vertices of B are referred to as black. We denote by $G_t = \langle W \uplus B_t, E_t \rangle$ the bipartite graph after the t 'th black vertex has arrived. The graph G_t is constructed online in the following manner:

- we start with $G_0 = \langle W \uplus B_0, E_0 \rangle = \langle W \uplus \emptyset, \emptyset \rangle$;
- in turn t a new vertex $b_t \in B$ together with all its incident edges $E(b_t)$ is revealed and G_t is defined as:

$$\begin{cases} E_t = E_{t-1} \cup E(b_t), \\ B_t = B_{t-1} \cup \{b_t\}; \end{cases}$$

- we finish when $t = |W|$, i.e., the number of black vertices is equal to the number of white vertices.

The goal of our algorithm is to compute for each G_t the maximum size matching M_t . The final graph $G_{|W|}$ which is obtained in this process will be denoted by $G = \langle W \uplus B, E \rangle$. We denote $n = |W| = |B|$ and $m = |E|$.

For every $t \in [n]$, we add orientation to edges of the graph G_t . This orientation is induced by matching M_t : the matched edges are oriented towards black vertices, while the unmatched edges are oriented towards white vertices. When a new vertex b_t arrives, we get an intermediate orientation $G_t^{\text{int}} = (E_t^{\text{int}}, B_t)$,

where the edges of b_t are oriented towards its neighbors, and the rest of the edges is oriented according to M_{t-1} . Note that G_t^{int} and G_{t-1} differ only by one vertex b_t . Any simple directed path in G_t^{int} from b_t to some unmatched white vertex is an augmenting path. In turn t , if b_t can be matched, the edges of G_t^{int} are reoriented along some augmenting path chosen by the algorithm, and the resulting orientation is G_t . The unmatched white vertices are called *seeds*. We denote the set of seeds in turn t as

$$S_t = \{w \in W : wb \notin M_t \text{ for any } b \in B\}.$$

So in turn t the augmenting paths in G_t^{int} are the directed paths from b_t to some $s \in S_{t-1}$. We represent a path as an ordered sequence of vertices. We use the notation $v \xrightarrow{\pi} v'$ to denote that a (directed) path π starts in v and ends in v' . We use the notation $v \in \pi$ and $\rho \subseteq \pi$ to state that a vertex v is a member of the sequence representing π and that a path ρ is a consecutive subsequence of π , respectively.

As mentioned in Section I, we study in this paper a greedy approach. The idea, explained formally further, is as follows. We associate with every white vertex (server) an attribute called *rank*. It denotes how many times the server was re-allocated, i.e., how many times was the corresponding white vertex a member of an augmenting path. We wish to minimize the maximum rank over all white vertices. A *rank of a path* π is defined to be equal to the maximum rank among the white vertices on π . In every turn t , if b_t can be matched, our algorithm chooses an augmenting path $b_t \xrightarrow{\pi_t} s \in S_{t-1}$ such that every suffix of π_t is of minimum rank.

We now introduce this idea formally. We define the rank function on W for every turn t . We set $\text{rank}_0(w) = 0$ for every $w \in W$ (it is always the case that $\text{rank}_t(w) = 0$ iff $w \in S_t$, i.e., w is a seed in turn t). The rank_t function is obtained from rank_{t-1} by increasing by one the rank of every white vertex $w \in \pi_t$:

$$\text{rank}_t(w) = \begin{cases} \text{rank}_{t-1}(w) & \text{if } w \notin \pi_t, \\ \text{rank}_{t-1}(w) + 1 & \text{if } w \in \pi_t. \end{cases}$$

We define the rank of a path π in G_t (or in G_{t+1}^{int}) as

$$\text{rank}_t(\pi) = \max_{w \in \pi} \text{rank}_t(w).$$

In addition to the rank function on the white vertices we introduce a *tier* function on the black vertices. The *tier* of a black vertex is the minimal rank of a directed path in G_t^{int} leading to a seed. For any $t \in [n]$ and $b \in B_t$ we define

$$\text{tier}_t(b) = \begin{cases} \infty & \text{if there is no path from } b \text{ to } S_t, \\ \min_{s \in S_t, (b \xrightarrow{\pi} s) \subseteq G_t} \text{rank}_t(\pi) & \text{otherwise.} \end{cases}$$

Observe that for every white vertex w there is at most one outgoing edge in every G_t (i.e., the matched edge) and denote the black vertex matched with w by b_w^t (we skip the upper index when it is clear from the context). Each edge wb_w^t in G_t from a white vertex w to the black vertex b_w^t is *tired* in G_t . Moreover, edge bw is *tiered* if the tier of b is not smaller

than the tier of b_w^t and the rank of w . The set of all tiered edges at the beginning of turn t is denoted by $E_t^* \subseteq E_t^{\text{int}}$,

$$E_t^* = \{bw \in E_t^{\text{int}} : b \in B, w \in W \text{ and} \\ \text{tier}_{t-1}(b) \geq \max\{\text{tier}_{t-1}(b_w), \text{rank}_{t-1}(w)\}\} \\ \cup \{wb_w^{t-1} \in E_t^{\text{int}} : w \in W\}.$$

This definition is extended to paths by saying that a path is *tiered* if it is a simple path that consists entirely of tiered edges.

In each round t , if possible, our algorithm chooses in G_t^{int} a tiered path $b_t \xrightarrow{\pi_t} s \in S_{t-1}$. The definition of a tiered path enforces that each suffix of π_t is optimal in a sense that it minimizes the rank. The algorithm then changes the orientation of the edges on π_t and increases the ranks of the white vertices. We refer to any such algorithm as a *tiered* algorithm. In the next sections we first show several properties of all tiered algorithms, and only after that we show how exactly our algorithm efficiently chooses the paths.

To conclude this section, we introduce a structure that will become useful in many of our proofs later. This structure also explains why the unmatched white vertices are called seeds. A *seeded tree* τ is any directed tree rooted at some seed $s \in S_t$. Any forest of $|S_t|$ vertex-disjoint seeded trees that spans all vertices of G that can reach a seed is called *seeded*. We call a seeded forest *tiered* if it consists entirely of tiered edges. A *seeded path* is any simple path to a seed.

Observe that not all vertices are contained in some seeded tree. The algorithm may produce orientations in which some vertices cannot reach any seed. We essentially ignore these vertices, as there is no need for the algorithm to visit them. Once a vertex cannot reach any seed, it will never be able to reach a seed again. Indeed, consider the set of vertices reachable from such a vertex. This set contains no seeds. Moreover, all the edges between this set and the rest of vertices are directed towards this set. Hence, no alternating path can reorient the edges between the set and the remaining part of the graph. We refer to vertices that cannot reach a seed as *dead*, the remaining vertices are called *alive*. As a consequence of the above argument, it only makes sense to consider the turns in which b_t can be matched. In the remaining turns neither the alive part of the graph changes, nor do the ranks. Thus, from now on, we completely ignore the black vertices that cannot be matched and the turn index t corresponds to the turns where the augmenting path can be found.

III. RANKS AND TIERS

In this section we present basic properties of ranks and tiers when they are maintained by tiered algorithms. These results will form the basis for the following sections. We start this section by proving that tiers of vertices do not decrease. Next, we prove that vertices with high tiers are far from the seeds. To be more precise, we show that for a vertex b of finite tier, the length of the shortest path from b to a seed is at least its tier. Let us start with the following observations.

Observation III.1. The following facts hold for $t \in [n]$:

- 1) There exists a tiered path from any alive vertex $b \in B_t$ to a seed in S_t .
- 2) For any tiered path π from $b \in B_t$ to a seed $s \in S_t$ it holds that $\text{tier}_t(b) = \text{rank}_t(\pi)$.
- 3) There exists a seeded tiered forest in G_t .
- 4) If there is a non-simple directed path from b to w in G_t , then there is also a simple directed path from b to w with the same or smaller rank.

We are ready to prove the following important lemma.

Lemma III.2. For any vertex $b \in B_t$ it holds that $\text{tier}_t(b)$ is a non-decreasing function of t .

Proof. Fix some seeded tiered forest ψ in G_t , which exists by Observation III.1.1. First assume that b is alive, so $b \in B(\psi)$. Let $\text{dist}_\psi(b)$ denote the black distance from b to a seed in ψ , i.e., the number of black vertices on a path from b to the corresponding seed (not including b). We proceed by induction on $\text{dist}_\psi(b)$. The vertices with $\text{dist}_\psi(b) = 0$ can directly reach seeds, which means that $\text{tier}_t(b) = 0$. A non-seed cannot turn into a seed, so $\text{tier}_{t-1}(b) = 0$ for such vertices.

Let b be an arbitrary alive vertex of B_t . Let $\pi \subseteq \psi$ be (the only) tiered path in ψ connecting b with a seed. Let w be a direct successor of b on π and b' be a black successor of w on π (the case when w is a seed was already covered). Since $\text{dist}_\psi(b') < \text{dist}_\psi(b)$, by induction we have $\text{tier}_t(b') \geq \text{tier}_{t-1}(b')$. By Observation III.1.2, since π is tiered, it holds that $\text{tier}_t(b) = \text{rank}_t(\pi) = \max\{\text{rank}_t(w), \text{tier}_t(b')\}$. Denote by ρ the path along which the edges were reversed in turn t . We distinguish three cases shown in Figure III.1.




Fig. III.1. Three cases depending the location of vertex b with regard to the augmenting path ρ .

If $w \notin \rho$, then bwb' is a directed path in G_{t-1} , so

$$\text{tier}_t(b) \geq \max\{\text{rank}_t(w), \text{tier}_t(b')\} \\ \geq \max\{\text{rank}_{t-1}(w), \text{tier}_{t-1}(b')\} \\ \geq \text{tier}_{t-1}(b).$$

The last inequality follows from Observation III.1.4.

If $b'wb \subseteq \rho$ (i.e., the path bwb' changed its orientation), since ρ is a tiered path in turn $t-1$, it holds that $\text{tier}_{t-1}(b') \geq \text{tier}_{t-1}(b)$. Therefore

$$\text{tier}_t(b) \geq \text{tier}_t(b') \geq \text{tier}_{t-1}(b') \geq \text{tier}_{t-1}(b).$$

Finally, it might be that $b'wb'' \subseteq \rho$ where $b'' \in B_t, b'' \neq b$. That means that there is a path with prefix bwb'' to a seed

$s \in S_{t-1}$ in G_{t-1} . On the other hand the smallest rank of the path from b to a seed in G_{t-1} equals $\text{tier}_{t-1}(b)$, so

$$\max\{\text{rank}_{t-1}(w), \text{tier}_{t-1}(b'')\} \geq \text{tier}_{t-1}(b). \quad (1)$$

Because b' lies on the tiered path (the fragment of tiered forest ψ) from b to a seed in G_t and b'' lies on the tiered path (the fragment of augmenting path ρ) from b' to a seed in G_{t-1} we obtain that

$$\text{tier}_t(b) \geq \text{tier}_t(b') \geq \text{tier}_{t-1}(b') \geq \text{tier}_{t-1}(b''). \quad (2)$$

Taking together (1) and (2) we obtain that

$$\text{tier}_t(b) \geq \max\{\text{rank}_t(w), \text{tier}_{t-1}(b'')\} \geq \text{tier}_{t-1}(b).$$

The case when b is a dead vertex, so its tier becomes or remains infinite, completes the proof. \square

The next lemma shows, that high ranks are followed by high tiers. It directly implies, that vertices of high tiers are far from the seeds.

Lemma III.3. *For any white vertex $w \in W \setminus S_t$, we have $\text{tier}_t(b_w^t) + 1 \geq \text{rank}_t(w)$.*

Proof. Assume that $\text{rank}_t(w) > 0$ (i.e., $w \notin S_t$) and let \tilde{t} be the last time the rank of w changed, i.e.,

$$\tilde{t} = \max\{i \in [t] : \text{rank}_{i-1}(w) < \text{rank}_i(w)\}.$$

Then, since the augmenting path $\pi_{\tilde{t}}$ was tiered, it holds that $\text{tier}_{\tilde{t}-1}(b_w^{\tilde{t}}) \geq \text{rank}_{\tilde{t}-1}(w)$. However, by Lemma III.2, the tier does not decrease, so

$$\begin{aligned} \text{tier}_t(b_w^t) + 1 &\geq \text{tier}_{\tilde{t}-1}(b_w^{\tilde{t}}) + 1 \geq \\ &\geq \text{rank}_{\tilde{t}-1}(w) + 1 = \text{rank}_t(w). \end{aligned}$$

\square

Corollary III.4. Let $B_t \ni b \xrightarrow{\pi_t} s \in S_t$. Every white vertex $w \in \pi$ of rank $\text{rank}_t(w) \geq 1$ is followed by some vertex w' of rank $\text{rank}_t(w') \geq \text{rank}_t(w) - 1$.

Proof. By Lemma III.3 $\text{tier}_t(b_w) \geq \text{rank}_t(w) - 1$. By definition of tier, any path (including π) from b_w to any seed contains a white vertex w' such that $\text{rank}_t(w') \geq \text{tier}_t(b_w)$. \square

Corollary III.5. The length of the shortest path from $b \in B_t$ to a seed is at least $2 \text{tier}_t(b) + 1$.

IV. UPPER BOUND ON WHITE VERTEX USAGE

In this section we prove the main result of the paper, which gives an $O(\sqrt{n})$ upper bound on the ranks of the vertices generated by any tiered algorithm. The proof proceeds in two stages. Lemma IV.1 captures a certain reversed behavior of a tiered algorithm crucial for our proof, pictured by Figure IV.1. Figure IV.1 presents the state of the graph after (to the left: graph G_t) and before (to the right: graph G_{t-1}) the augmentation step. Assume that graph G_t contains seeded paths $\mu_1 \dots \mu_l$. On the picture to the left, the ranks of these paths are shown by the heights of the corresponding bars, sorted by rank. Lemma IV.1 states, that before the

augmentation step (the right side of the picture), each path μ_i had its counterpart λ_i , plus there was one more path ρ . The interesting part is that the heights of the bars (the ranks of the corresponding paths) decrease at most by one, and the newly added bar (the one corresponding to ρ) dominates the bars that decrease.

Hence, the line inclined at the angle of 45 degrees through the right top point of the bar representing λ_l does not drop: whenever the bars reduce their height by one, an additional bar appears to support the line at its initial position. The idea




Fig. IV.1. The bars corresponding to the paths after and before the augmentation step.

is that for any white vertex whose rank is r at some point, there is a turn when a seeded path of rank at least r exists. So, if the algorithm produces a vertex of rank r , we start with one bar of height r and reverse the steps of the algorithm. We reach a point when there is a linear in r number of vertex disjoint paths whose height is also linear in r . The next lemmas formally prove our result.

It is worth mentioning that in the following reasoning it is enough to consider alive vertices. Dead vertices do not affect the definition of rank_t and tier_t and behavior of the algorithm.

Lemma IV.1. *Let $\mu_1, \mu_2, \dots, \mu_l$ be a sequence of vertex-disjoint seeded paths in G_t . Also, let $b_t \xrightarrow{\pi_t} s$ be the augmenting path in turn t , i.e., in G_t^{int} . Then, there exist vertex-disjoint seeded paths $\lambda_1, \dots, \lambda_l$ and ρ in G_{t-1} such that*

$$\text{rank}_{t-1}(\pi) \leq \text{rank}_{t-1}(\rho), \quad (3)$$

$$\text{rank}_t(\mu_i) \leq \text{rank}_{t-1}(\lambda_i) \quad (4)$$

$$\text{for } \text{rank}_t(\mu_i) > \text{rank}_{t-1}(\pi) + 1,$$

$$\text{rank}_t(\mu_i) \leq \text{rank}_{t-1}(\lambda_i) + 1 \quad (5)$$

otherwise.

Proof. Let w_i be the highest-ranked white vertex of μ_i that is the closest to the seed. Without loss of generality we can assume that μ_i starts with w_i . Then, for $i \in [l]$, directed path μ_i is of the form $W \ni w_i \xrightarrow{\mu_i} s_i \in W$ where $s_i \in S_t$ is a seed.

Now, we define a 0/1-flow network F where $V(F) = V(G_t^{\text{int}}) \cup \{\zeta, \tau\}$ and

$$E(F) = E(G_t^{\text{int}}) \cup \{\zeta b_t, \zeta w_1, \dots, \zeta w_l\} \cup \{s\tau, s_1\tau, \dots, s_l\tau\}$$

where ζ and τ are artificially added source and sink respectively (see Fig. IV.2).




Fig. IV.2. Example of a flow network F , where solid arcs form the flow from source ζ to sink τ with the total value 4.

Because $b_t \xrightarrow{\pi_t} s$ is an augmenting path expanding matching in graph G_t^{int} , we have that $\zeta b_t \xrightarrow{\pi_t} s \tau$ is a path which defines a flow f_1 in F with the total flow value equal to 1. Moreover, we obtain that the residual graph F^{f_1} in restriction to $V(G_t)$ has the same orientation as G_t . Because μ_1, \dots, μ_l are pairwise vertex-disjoint directed paths in G_t , we have that paths μ_i of the form $\zeta w_i \xrightarrow{\mu_i} s_i \tau$ are pairwise edge-disjoint directed paths in residual graph F^{f_1} . Consecutive application of the augmenting paths μ_1, \dots, μ_l expands the flow f_1 to a flow f_{l+1} with the total value of $l+1$.

The flow f_{l+1} determines $l+1$ edge-disjoint paths in F from ζ to τ as f_{l+1} is a $0/1$ -flow. Because b_t, w_1, \dots, w_l are all neighbors of ζ and $s_1, \dots, s_{l+1} = s$ are all neighbors of τ , we obtain $l+1$ paths of the form $\zeta b_t w_{l+1} \xrightarrow{\rho} s_{\sigma(l+1)} \tau$ and $\zeta w_i \xrightarrow{\lambda_i} s_{\sigma(i)} \tau$ for $i \in [l]$ where σ is some permutation of $[l+1]$ and $\lambda_1, \dots, \lambda_l, \rho$ are some edge-disjoint directed paths in $G_t^{\text{int}} \subseteq F$. Because w_1, \dots, w_l, b_t are different and s_1, \dots, s_l, s_{l+1} are also different and each vertex in G_{t-1} has out-degree or in-degree equals at most 1, paths $\lambda_1, \dots, \lambda_l, \rho$ are also vertex-disjoint.

To see (3), it is enough to mention that path π was chosen in turn t so it was the smallest rank seeded path from b_t . Because the augmenting path π in turn t uses white vertices with rank at most $\text{rank}_{t-1}(\pi)$, lines (4)-(5) are obvious. \square

The next lemma presents the core of our result. Imagine that for some white vertex w its rank was raised to $\text{rank}_{\tilde{t}}(w)$ by some turn \tilde{t} . The lemma states that we will find a turn before turn \tilde{t} , when a collection of vertex-disjoint paths existed, such that the sum of their ranks was at least quadratic with respect to $\text{rank}_{\tilde{t}}(w)$. The proof heavily bases on Lemma IV.1, which allows us to reverse the steps of the algorithm.

We begin the reversing process in turn t' , where a path ρ exists with rank at least $\text{rank}_{\tilde{t}}(w) - 1$. Starting from a one element collection $\{\rho\}$, we iteratively construct a collection of paths using Lemma IV.1. From the collection in step t we obtain a collection in step $t-1$. There are two important invariants in this reversing process which imply the lemma:

- minimum rank in the collection decreases at most by one

- if the minimum rank in the collection decreases, the size of the collection grows by one
- there is a moment t_0 when all ranks in the collection are equal to 0

Somewhere before t_0 there is a moment, where there are sufficiently many paths of sufficiently high ranks, as illustrated in Figure IV.3. Again, each bar represents a path and the heights of the bars are the ranks of the corresponding paths. Let us now move on to formalizing this idea.




Fig. IV.3. In time t there is a collection of $\lfloor \text{rank}_{\tilde{t}}(w)/2 \rfloor$ paths of rank at least $\lceil \text{rank}_{\tilde{t}}(w)/2 \rceil$.

Lemma IV.2. For any turn $\tilde{t} \in [n]$ and any white vertex $w \in W_{\tilde{t}}$ there is $t \leq \tilde{t}$ and j pairwise vertex-disjoint seeded simple paths P_1, \dots, P_j in G_t such that

$$\left\lfloor \frac{1}{4} \text{rank}_{\tilde{t}}(w)^2 \right\rfloor \leq \text{rank}_t(P_1) + \dots + \text{rank}_t(P_j).$$

Proof. Let $t' \leq \tilde{t}$ be the round when the rank of white vertex w was changed the last time before round \tilde{t} . Now, we define sequence $B_{t'-1}, \dots, B_0$, where B_k is a collection of edge-disjoint simple seeded paths in G_k for $k = 0, \dots, t'$. From the definition of t' we know that $\text{rank}_{\tilde{t}}(w) = \text{rank}_{t'}(w)$ and $\text{rank}_{t'-1}(\pi_{t'}) \geq \text{rank}_{t'-1}(w)$. If we remove the vertex $b_{t'}$ from $\pi_{t'}$, we obtain a simple path ρ in $G_{t'-1}$ with rank

$$\begin{aligned} \text{rank}_{t'-1}(\rho) &= \text{rank}_{t'-1}(\pi_{t'}) \geq \text{rank}_{t'-1}(w) = \\ &= \text{rank}_{t'}(w) - 1 = \text{rank}_{\tilde{t}}(w) - 1. \end{aligned}$$

Thus, we define $B_{t'-1} = \{\rho\}$. Now, let us assume that $B_k = \{\mu_1, \dots, \mu_l\}$ is defined for some $k \in \{1, \dots, t'-1\}$. Again, by applying Lemma IV.1 to G_k , a sequence of paths μ_1, \dots, μ_l , and augmenting path π_k , we obtain simple seeded paths $\lambda_1, \dots, \lambda_l, \bar{\rho}$ in G_{k-1} satisfying conditions (1)-(3). Then we define B_{k-1} as

$$B_{k-1} = \begin{cases} \{\lambda_1, \dots, \lambda_l, \bar{\rho}\} & \exists_{i \in [l]} \text{rank}_k(\mu_i) > \text{rank}_{k-1}(\lambda_i), \\ \{\lambda_1, \dots, \lambda_l\} & \text{otherwise (i.e., } \forall_{i \in [l]} \text{rank}_k(\mu_i) = \text{rank}_{k-1}(\lambda_i).) \end{cases}$$

If $\text{rank}_k(\mu_i) > \text{rank}_{k-1}(\lambda_i)$ for some $i \in [l]$ then $\text{rank}_k(\mu_i) \leq \text{rank}_{k-1}(\pi_k) + 1 \leq \text{rank}_{k-1}(\bar{\rho}) + 1$, by conditions (1)-(3) from Lemma IV.1. Again, applying inequalities (1)-(3) from Lemma IV.1, we obtain that

$$\min_{\lambda \in B_{k-1}} \text{rank}_{k-1}(\lambda) \geq \begin{cases} \min_{\mu \in B_k} \text{rank}_k(\mu) - 1 & \text{if } |B_{k-1}| - |B_k| = 1, \\ \min_{\mu \in B_k} \text{rank}_k(\mu) & \text{otherwise (i.e., } |B_{k-1}| - |B_k| = 0). \end{cases}$$

Multiple use of the above inequality gives us

$$\begin{aligned} \min_{\lambda \in B_k} \text{rank}_k(\lambda) &\geq \min_{\mu \in B_{t'-1}} \text{rank}_{t'-1}(\mu) - (|B_k| - 1) \\ &= \text{rank}_{\bar{t}}(w) - |B_k|, \end{aligned} \quad (6)$$

for any chosen $k = 0, \dots, t' - 1$. The formula (6) implies that family B_0 is the size of at least $\text{rank}_{\bar{t}}(w)$ by $\text{rank}_0(\lambda) = 0$ for all $\lambda \in B_0$. Thus, we can choose $t \in [t']$ for which $|B_t| = \lfloor \frac{1}{2} \text{rank}_{\bar{t}}(w) \rfloor$ and by (6) we have $\min_{\lambda \in B_t} \text{rank}_t(\lambda) \geq \lceil \frac{1}{2} \text{rank}_{\bar{t}}(w) \rceil$. Finally, we obtain the family $B_t = \{P_1, \dots, P_j\}$ of size $j = |B_t| = \lfloor \frac{1}{2} \text{rank}_{\bar{t}}(w) \rfloor$ where each $P_i \in B_t$ has the rank at least $\lceil \frac{1}{2} \text{rank}_{\bar{t}}(w) \rceil$. This together gives a desired lower bound on the sum of the ranks of P_i :

$$\begin{aligned} \left[\frac{1}{4} \text{rank}_{\bar{t}}(w)^2 \right] &\leq \left[\frac{1}{2} \text{rank}_{\bar{t}}(w) \right] \cdot \left[\frac{1}{2} \text{rank}_{\bar{t}}(w) \right] \\ &\leq |B_t| \cdot \min_{P_i \in B_t} \text{rank}_t(P_i) \\ &\leq \text{rank}_t(P_1) + \dots + \text{rank}_t(P_j). \end{aligned} \quad \square$$

Now, we are ready to prove the main theorem:

Theorem IV.3. *For any tiered algorithm it holds that $\text{rank}_n(w) < \sqrt{2n}$ for every $w \in W$.*

Proof. By Lemma IV.2 there is round $t \in [n]$ in which there is a collection of j pairwise vertex-disjoint directed simple paths $P_1, \dots, P_j \in G_t$ for which the sum of their ranks is $\lfloor \frac{1}{4} \text{rank}_n(w)^2 \rfloor$. This gives that P_1, \dots, P_j are paths of total length greater than $\frac{1}{2} \text{rank}_n(w)^2$ by Corollary III.5. On the other hand, graph G_t has at most n vertices, so $\text{rank}_n(w) < \sqrt{2n}$. \square

V. THE ALGORITHM

In the previous sections we showed that any tiered algorithm rematches a white vertex at most $O(\sqrt{n})$ times. This shows that the sum of the lengths of all used augmenting paths is bounded by $O(n\sqrt{n})$. Unfortunately the algorithm not only needs to apply the augmenting paths, but also needs to find them, so the total running time might be much higher than this bound. In this section we show an incremental algorithm finding maximal matching with a total running time of $O(m\sqrt{n})$, which is equal to the running time of the Hopcroft–Karp algorithm.

In essence, we observe that we can pay for the search by increasing the ranks of the visited white vertices, and the bound of $O(\sqrt{n})$ for the white vertex usage still holds.

Algorithm *Match* proposed here proceeds in the following way (see Algorithm 1 for the pseudocode). At the initialization time, *Match* receives G_0 as an input, and sets $w.\text{rank}^* = 0$ for every white vertex $w \in W$. We denote the ranks produced by *Match* in turn t as $\text{rank}_t^*(w)$. So, due to the initialization step, for every white vertex $w \in W$ we have $\text{rank}_0^*(w) = 0$. The tier function for every $t \in [n]$ is defined, as before, based on the rank function, i.e., for any $b \in B_t$

$$\text{tier}_t^*(b) = \begin{cases} \infty & \text{if there is no path from } b \text{ to } S_t, \\ \min_{s \in S_t, b \xrightarrow{\pi} s} \text{rank}_t^*(\pi) & \text{otherwise.} \end{cases}$$

In turn t , the algorithm searches for the tiered augmenting path π_t starting at b_t . It proceeds recursively. The recursive function $\text{SEARCH}(w)$ is called for a white vertex w , initially equal to the white neighbor of b_t with the smallest rank. If w is a seed ($\text{UNMATCHED}(w)$ returns `true`) then w increases its rank and, since a seeded path is found, $\text{SEARCH}(w)$ returns `true`. Otherwise the recursive step is determined by the black vertex b_w^{t-1} (matched with w). The algorithm (using function $\text{SMALLESTNEIGHBORUNMATCHEDTO}(b_w^{t-1})$) picks the white neighbor w' of b_w^{t-1} with the minimum rank, i.e.,

$$w'.\text{rank}^* = \min_{w'' : b_w^{t-1}w'' \in E_t^{\text{int}}} w''.\text{rank}^*.$$

The value of $w'.\text{rank}^*$ is a lower bound for $\text{tier}_{t-1}^*(b_w^{t-1})$. A successful recursive call on w' recursively finds a directed path to a seed with rank bounded from above by $w'.\text{rank}^*$. Hence, if the call is successful, $\text{tier}_{t-1}^*(b_w^{t-1}) = w'.\text{rank}^*$. This certifies that both edges wb_w^{t-1} and $b_w^{t-1}w'$ are tiered and so the algorithm found a tiered path. In such case the recursion terminates returning `true` and the ranks are increased along the tiered path that was found. Each time a rank of a white vertex w is increased, the function $\text{UPDATEINFORMATION}(b, w)$ is called in order to update $w.\text{rank}^*$ for all the black vertices b that can directly reach w . An unsuccessful recursive call on w' certifies, that $\text{tier}_{t-1}^*(b_w^{t-1}) > w'.\text{rank}^*$. On such an event, the algorithm increases $w'.\text{rank}^*$ and continues picking the neighbors of b_w^{t-1} of minimum rank and calling the recursion on them. This process finishes when either the call is successful or the only available ranks reach the bound of $\sqrt{2n}$. In the latter case the algorithm located a dead region, which never will be entered again. Note that every time a white vertex is visited by the algorithm, its rank increases. We soon show, that unless a vertex is dead, its rank is bounded by $\sqrt{2n}$.

As for the running time, note that the function $\text{UPDATEINFORMATION}(b, w)$ can be handled in constant time for every neighbor b pointing at w . For example each black vertex can keep a list of lists of white vertices. This allows to update the information in this recursive call in the time bounded by the degree of w . At the same time the rank of w is increased, and this happens at most $O(\sqrt{n})$ times for any white vertex. The total running time of $O(m\sqrt{n})$ follows.

Consider an arbitrarily chosen vertex w . We distinguish two phases of the algorithm *Match* in turn t , encoded as $\text{MATCH}(b_t)$ in Algorithm 1. The first phase consists of a sequence (which can be empty) of calls to $\text{SEARCH}(w)$ which return `false`. Each such call does not rematch w but increases $w.\text{rank}^*$. The second phase consists of at most one call $\text{SEARCH}(w)$ with returns `true`. In such call, $\text{SEARCH}(w)$ rematches vertex w and increases $w.\text{rank}^*$ as well. We formulate the following lemma, which can be proven by induction on the number of calls to $\text{SEARCH}(\cdot)$. The lemma describes the difference between the above mentioned two phases.

Lemma V.1. *Consider any recursive call $\text{SEARCH}(w)$ originating from a call to $\text{MATCH}(b_t)$ for any $w \in W$. There is a path p from w to some seed with ranks on white vertices not greater than $w.\text{rank}^*$ if and only if $\text{SEARCH}(w)$ returns `true`*

Algorithm 1 Algorithm *Match*

```
1: procedure MATCH( $b_t$ )
2:    $w' \leftarrow$  SMALLESTNEIGHBORUNMATCHEDTO( $b_t$ )
3:   while  $w'.rank^* < \sqrt{2n}$  do
4:     if SEARCH( $w'$ ) then
5:        $match \leftarrow match \cup \{(b_t, w')\}$ 
6:     return true
7:    $w' \leftarrow$  SMALLESTNEIGHBORUNMATCHEDTO( $b_t$ )
8:   return false

9: procedure SEARCH( $w$ )
10:   $w.rank^* \leftarrow w.rank^* + 1$ 
11:  for each  $b$  such that  $bw \in E_t^{int}$  do
12:    UPDATEINFORMATION( $b, w$ )
13:  if UNMATCHED( $w$ ) then
14:    return true
15:   $w' \leftarrow$  SMALLESTNEIGHBORUNMATCHEDTO( $b_w$ )
16:  while  $w'.rank^* < w.rank^*$  do
17:    if SEARCH( $w'$ ) then
18:       $match \leftarrow match \setminus \{(w, b_w)\} \cup \{(b_w, w')\}$ 
19:    return true
20:   $w' \leftarrow$  SMALLESTNEIGHBORUNMATCHEDTO( $b_w$ )
21:  return false
```

and rematches edges along some path π from w to some seed with ranks on white vertices smaller or equal to $w.rank^*$.

One of the consequences of the above lemma is that during the first phase, although $w.rank^*$ grows, it does not become greater than $\text{tier}_{t-1}^*(b_w^{t-1})$. Let us define $\text{rank}_{t-1}^{\otimes}(w)$ as a value $w.rank^*$ immediately after the first phase. It holds that

$$\begin{aligned} & \text{if } \text{rank}_{t-1}^{\otimes}(w) > \text{rank}_{t-1}^*(w) \\ & \text{then } \text{rank}_{t-1}^{\otimes}(w) \leq \text{tier}_{t-1}^*(b_w^{t-1}) \end{aligned} \quad (7)$$

and so we obtain the following:

Observation V.2. For any $t \in [n]$ and any path $w \xrightarrow{\pi} s$, $w \in W, s \in S_{t-1}$, it holds that $\text{rank}_{t-1}^{\otimes}(\pi) = \text{rank}_{t-1}^*(\pi)$.

Now if b_t is dead, procedure MATCH(b_t) tries to match it anyway. Due to Observation V.2, all vertices visited during such a call to MATCH(b_t) incur only the first phase, and as a consequence the ranks of the paths do not change. This allows us to simplify the description and omit the rounds when dead vertices appear.

With respect to the second phase, we can observe that during that phase the algorithm increases by one ranks of white vertices on the augmenting path π_t , i.e.,

$$\text{rank}_t^*(w) = \begin{cases} \text{rank}_{t-1}^{\otimes}(w) & \text{if } w \notin \pi_t, \\ \text{rank}_{t-1}^{\otimes}(w) + 1 & \text{if } w \in \pi_t. \end{cases} \quad (8)$$

By Observation V.2 and equation (8) very similar arguments as for Lemma III.2 and Lemma III.3 prove the corresponding lemmas:

Lemma V.3. For any vertex $b \in B_t$ it holds that $\text{tier}_t^*(b)$ is a non-decreasing function of t .

Lemma V.4. For any white vertex $w \in W \setminus S_t$, we have $\text{tier}_t^*(b_w) + 1 \geq \text{rank}_t^*(w)$.

The next two corollaries follow.

Corollary V.5. Let $b \xrightarrow{\pi} s$, for a black vertex b , be a path that ends at a seed s . Every white vertex $w \in \pi$ of rank $\text{rank}_t^*(w) \geq 1$ is followed by some vertex w' of rank $\text{rank}_t^*(w') \geq \text{rank}_t^*(w) - 1$.

Corollary V.6. The length of the shortest path from a black vertex b to a seed is at least $\text{tier}_t^*(b)$.

We move on to the two lemmas crucial for our approach.

Lemma V.7. Let $\mu_1, \mu_2, \dots, \mu_l$ be a sequence of vertex-disjoint seeded simple paths to distinct seeds in G_t . Also, let $b_t \xrightarrow{\pi_t} s$ be the augmenting path in turn t , i.e., in G_{t-1}^{int} . Then, there exist vertex-disjoint simple seeded paths $\lambda_1, \dots, \lambda_l$ and ρ in G_{t-1} to distinct seeds such that

$$\text{rank}_{t-1}^*(\pi_t) \leq \text{rank}_{t-1}^*(\rho), \quad (9)$$

$$\text{rank}_t^*(\mu_i) \leq \text{rank}_{t-1}^*(\lambda_i) \quad (10)$$

$$\text{for } \text{rank}_t^*(\mu_i) > \text{rank}_{t-1}^*(\pi) + 1,$$

$$\text{rank}_t^*(\mu_i) \leq \text{rank}_{t-1}^*(\lambda_i) + 1 \quad (11)$$

otherwise.

Idea of proof. The analogous argument as in proof of Lemma IV.1 implies that the paths $\lambda_1, \dots, \lambda_l$ and ρ as described exist in G_t^{int} . Due to Observation V.2 and the equation (8) we obtain the desired inequalities. \square

Now, the proof of Lemma IV.2 directly translates into an argument for the following lemma:

Lemma V.8. For any turn $\tilde{t} \in [n]$ and any white vertex $w \in W_{\tilde{t}}$ there is $t \leq \tilde{t}$ and j pairwise vertex-disjoint seeded simple paths P_1, \dots, P_j in G_t such that

$$\left\lfloor \frac{1}{4} \text{rank}_{\tilde{t}}^*(w)^2 \right\rfloor \leq \text{rank}_{\tilde{t}}^*(P_1) + \dots + \text{rank}_{\tilde{t}}^*(P_j).$$

As a consequence, we obtain a corollary analogous to Theorem IV.3.

Corollary V.9. For any tiered algorithm it holds that $\text{rank}_n^*(w) < \sqrt{2n}$ for every $w \in W$.

VI. APPLICATIONS

In this section we describe few applications of our result. First, we present a linear time incremental approximation algorithm which, for any constant ε , maintains a $(1 - \varepsilon)$ -approximation of maximum matching in a bipartite graph revealed online. Interestingly, if we settle for approximation, we do not need Theorem IV.3 for the algorithm to work. This algorithm implies a simple offline algorithm with a simple analysis, also not requiring Theorem IV.3. It shows that the real value of our result lies in the online nature of the problem.

Further, we reduce the decremental version of the problem to the incremental version, both in exact and approximate case. We finish the section by generalizing our result to the weighted setting.

A. Incremental approximation and offline analysis

The approximation algorithm is based on the following well known fact, that a matching for which all augmenting paths are long is a good approximation of an optimum matching in a bipartite graph.

Claim VI.1. *Let M^* be an optimum matching in a bipartite graph B , and M be a matching. If the length of the shortest path augmenting M is at least k , then $|M| \geq |M^*|(1 - \frac{2}{k})$.*

Combined with the above claim, our algorithm gives an efficient $(1 - \varepsilon)$ -approximation for incremental bipartite matching. We try to match a black vertex b_t as long as $\min_{w : b_t w \in E_t^{\text{int}}} \text{rank}(w) \leq k$ for some constant k . Once this value is higher than k , we leave b_t unmatched. In every turn, due to Lemma III.5, all augmenting paths starting at vertices we left out are of length at least k . Hence, in each turn we maintain a $(1 - \frac{2}{k})$ -approximation of the optimum matching. Moreover, since we took care that ranks are bounded by k , the total time used by the approximation algorithm is $O(km)$. We summarize this interesting observation in the following lemma. It is worth emphasizing that we do not need Theorem IV.3 for this argument: we impose low ranks ourselves here.

Lemma VI.2. *For any $\varepsilon > 0$, there exists an incremental algorithm that maintains $(1 - \varepsilon)$ -approximate maximum bipartite matching in $O(\varepsilon^{-1}m)$ total time.*

We use a similar idea in order to provide a simple analysis for the offline version of our algorithm. We apply the approximation algorithm with $2/\varepsilon = k = \sqrt{n}$. The running time of the approximation algorithm is, due to Lemma VI.2, $O(m\sqrt{n})$. We are left with at most $2|M^*|/k \in O(\sqrt{n})$ unmatched vertices, so we can afford to perform the full search for the augmenting paths. However, this simple analysis does not carry over if we need to match every black vertex at the moment it is presented.

B. Decremental Algorithm

Let us consider the online problem where at the beginning we are given a graph $G^d = \langle W^d \uplus B^d, E^d \rangle$, and the vertices from W are removed from G^d in an online manner. During this process we are asked to maintain the maximum size matching in G^d . More formally we consider graphs $G_t^d = (W_0^d \uplus B^d, E_t^d)$ constructed by the following scenario:

- we start with $G_0^d = G^d$,
- in turn t we remove vertex $w_t \in W$ along with all its adjacent edges:

$$\begin{cases} E_t^d = E_{t-1}^d \setminus \{w_t, b\} : b \in B^d, \\ W_t^d = W_{t-1}^d \setminus \{w_t\}. \end{cases}$$

- we finish when $t = |W_0^d|$, i.e., when all white vertices are removed.

Observe the important difference with respect to the incremental problem. Here, we are removing white vertices, whereas previously we were adding black vertices. This change in colors will be useful when reducing the decremental problem to the incremental one. Now we initialize the incremental algorithm with graph G^d , i.e., we set $G_0 = G^d$. When the vertex w_t should be removed from G_{t-1}^d we add a vertex b_t to G_{t-1} using the incremental algorithm. The only edge adjacent to b_t that we add to G_{t-1} is $b_t w_t$. Let G_t denote the sequence of graphs resulting from this incremental process. We observe the following:

Lemma VI.3. *Let M_t be the maximum size matching in G_t , and let M_t^d be the maximum size matching in G_t^d . Then $|M_t| - t = |M_t^d|$.*

Proof. Consider the matching $M_t' = M_t^d \cup \{b_i w_i : 1 \leq i \leq t\}$. By this construction we have $|M_t'| - t = |M_t^d|$. If M_t' is not a maximum size matching in G_t then there exists an augmenting path π with respect to M_t' . This path cannot pass via any of the edges in $\{b_i w_i : 1 \leq i \leq t\}$, because b_i has degree 1. Hence, π is an augmenting path for M_t^d in G_t^d what contradicts the maximality of M_t^d . Hence, $|M_t| = |M_t'| = |M_t^d| + t$. \square

Combining these results with the algorithm of Section V we obtain the following observation.

Corollary VI.4. *There exists a decremental algorithm for the maximum bipartite matching problem that works in $O(\sqrt{nm})$ total time.*

We will now prove that the same reduction works in the approximate case as well. In the incremental case we were using Claim VI.1. This fact cannot be used directly with the above reduction, because the sizes of the two matchings M_t and M_t^d differ by t . However, it is possible to prove the guarantee on approximation directly from the nonexistence of short augmenting paths.

Lemma VI.5. *Let M_t be a matching in G_t such the shortest augmenting path with respect to M_t has length k , for some $k \geq 1$, and let M_t^d be the maximum size matching in G_t^d . Then $|M_t| - t \geq \left(1 - \frac{2}{k-2}\right) |M_t^d|$.*

By using the above lemma together with Lemma VI.2 we obtain the following observation.

Corollary VI.6. *For any $\varepsilon > 0$, there exists a decremental algorithm that maintains $(1 - \varepsilon)$ -approximate maximum bipartite matching in $O(\varepsilon^{-1}m)$ total time.*

C. Weighted Algorithm

In this section we are going to recall the unfolded graph technique that was introduced in [13]. Let $G_w = (V_w, E_w)$ be a graph where edge weights are given by a function $w : E \rightarrow [1, W]$ for some natural number W . The *unfolded graph* $G = (V, E)$ of G_w is defined as follows:

$$V = \{v_i : v \in V \text{ and } i \in [1, W]\},$$

$$E = \{(u_i v_{W-i+1}) : uv \in E \text{ and } i \in [1, W]\}.$$

We have the following lemma.

Lemma VI.7 (Lemma 4.1 of [13]). *The weight of the maximum weight matching in G_w is equal to the size of the maximum matching in G .*

Using the above lemma we can compute the weight of the maximum matching in G_w online by maintaining the unfolded graph G . The unfolded graph has Wm edges and Wn vertices. Hence, the following corollaries are immediate.

Corollary VI.8. There exist incremental and decremental algorithms that maintain the weight of the maximum weight matching in G_w that work in $O(W^{3/2}\sqrt{nm})$ total time.

Corollary VI.9. There exist incremental and decremental algorithms that maintain $(1 - \epsilon)$ -approximation of the weight of the maximum weight matching in G_w that work in $O(W\epsilon^{-1}m)$ total time.

VII. CONCLUSIONS

In this paper we have introduced a new greedy framework for solving online maximum matching problem in bipartite graphs. We believe that the introduction of this framework will inspire further work on development of efficient online algorithms for other combinatorial problems like: maximum-flow, min-cost maximum-flow, stable matchings, matroid intersection or non-bipartite matchings. In particular, we would like to stress the following open problems.

- 1) Is it possible to extend our results to the online non-bipartite matching problem, where vertices of the graph are revealed online? So far only $O(nm)$ time exact solution is known.
- 2) Can our framework be adopted to the incremental max-flow problem? Here, even in the case where all edge capacities are equal to one, no dynamic solution is known for general graph. We know only some dynamic algorithms in the case of planar graphs [11].
- 3) We note that our model is different from the edge update scenarios studied in [12], [20], [18], [4], [17], [9], [2], and it would be interesting to see whether our techniques can be applied in this model to get fast incremental or decremental algorithms?
- 4) As noted in the introduction the lower bounds developed in [1] do not apply to our model, as they talk about edge updates. Can one show similar lower bounds for vertex-updates?

REFERENCES

- [1] A. ABBOD AND V. VASSILEVSKA WILLIAMS, *Popular conjectures imply strong lower bounds for dynamic problems*, ArXiv e-prints, (2014).
- [2] A. ANAND, S. BASWANA, M. GUPTA, AND S. SEN, *Maintaining Approximate Maximum Weighted Matching in Fully Dynamic Graphs*, in IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2012), D. D'Souza, T. Kavitha, and J. Radhakrishnan, eds., vol. 18 of Leibniz International Proceedings in Informatics (LIPIcs), Dagstuhl, Germany, 2012, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, pp. 257–266.
- [3] Y. AZAR, *On-line load balancing*, in Theoretical Computer Science, Springer, 1992, pp. 218–225.
- [4] S. BASWANA, M. GUPTA, AND S. SEN, *Fully dynamic maximal matching in $o(\log n)$ update time*, in Proceedings of the 2011 IEEE 52Nd Annual Symposium on Foundations of Computer Science, FOCS '11, Washington, DC, USA, 2011, IEEE Computer Society, pp. 383–392.
- [5] T. M. CHAN, *Dynamic subgraph connectivity with geometric applications*, in Proceedings of the Thiry-fourth Annual ACM Symposium on Theory of Computing, STOC '02, New York, NY, USA, 2002, ACM, pp. 7–13.
- [6] T. M. CHAN, M. PĂTRAȘCU, AND L. RODITTY, *Dynamic connectivity: Connecting to networks and geometry*, SIAM Journal on Computing, 40 (2011), pp. 333–349. See also FOCS'08, arXiv:0808.1128.
- [7] K. CHAUDHURI, C. DASKALAKIS, R. D. KLEINBERG, AND H. LIN, *Online bipartite perfect matching with augmentations*, in INFOCOM'09, 2009, pp. 1044–1052.
- [8] E. GROVE, M.-Y. KAO, P. KRISHNAN, AND J. VITTER, *Online perfect matching and mobile computing*, in Algorithms and Data Structures, S. Akl, F. Dehne, J.-R. Sack, and N. Santoro, eds., vol. 955 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 1995, pp. 194–205.
- [9] M. GUPTA AND R. PENG, *Fully dynamic $(1 + \epsilon)$ -approximate matchings*, 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, 0 (2013), pp. 548–557.
- [10] J. E. HOPCROFT AND R. M. KARP, *An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs*, SIAM Journal on Computing, 2 (1973), pp. 225–231.
- [11] G. F. ITALIANO, Y. NUSSBAUM, P. SANKOWSKI, AND C. WULFF-NILSEN, *Improved algorithms for min cut and max flow in undirected planar graphs*, in Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing, STOC '11, New York, NY, USA, 2011, ACM, pp. 313–322.
- [12] Z. IVKOVIĆ AND E. LLOYD, *Fully dynamic maintenance of vertex cover*, in Graph-Theoretic Concepts in Computer Science, J. Leeuwen, ed., vol. 790 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 1994, pp. 99–111.
- [13] M.-Y. KAO, T. W. LAM, W.-K. SUNG, AND H.-F. TING, *A decomposition theorem for maximum weight bipartite matchings with applications to evolutionary trees*, in Proceedings of the 7th Annual European Symposium on Algorithms, 1999, pp. 438–449.
- [14] R. M. KARP, U. V. VAZIRANI, AND V. V. VAZIRANI, *An optimal algorithm for on-line bipartite matching*, in STOC, H. Ortiz, ed., ACM, 1990, pp. 352–358.
- [15] A. MEHTA, A. SABERI, U. VAZIRANI, AND V. VAZIRANI, *Adwords and generalized on-line matching*, in Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on, Oct 2005, pp. 264–273.
- [16] V. MIRROKNI, personal communication.
- [17] O. NEIMAN AND S. SOLOMON, *Simple deterministic algorithms for fully dynamic maximal matching*, in Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC '13, New York, NY, USA, 2013, ACM, pp. 745–754.
- [18] K. ONAK AND R. RUBINFELD, *Property testing*, Springer-Verlag, Berlin, Heidelberg, 2010, ch. Dynamic Approximate Vertex Cover and Maximum Matching, pp. 341–345.
- [19] M. PATRASCU, *Towards polynomial lower bounds for dynamic problems*, in Proceedings of the Forty-second ACM Symposium on Theory of Computing, STOC '10, New York, NY, USA, 2010, ACM, pp. 603–610.
- [20] P. SANKOWSKI, *Faster dynamic matchings and vertex connectivity*, in Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '07, Philadelphia, PA, USA, 2007, Society for Industrial and Applied Mathematics, pp. 118–126.