A Theory of Alternating Paths and Blossoms, from the Perspective of Minimum Length

Vijay V. Vazirani
*
University of California, Irvine, vazirani@ics.uci.edu, https://www.ics.uci.edu/ vazirani/

The Micali-Vazirani (MV) algorithm for finding a maximum cardinality matching in general graphs, which was published in 1980, remains to this day the most efficient known algorithm for the problem. The current paper gives the first complete and correct proof of this algorithm.

The MV algorithm resorts to finding minimum length augmenting paths. However, such paths fail to satisfy an elementary property, called *breadth first search honesty* in this paper. In the absence of this property, an exponential time algorithm appears to be called for — just for finding one such path. On the other hand, the MV algorithm accomplishes this, and additional tasks, in linear time. The saving grace is the various "footholds" offered by the underlying structure, which the algorithm uses in order to perform its key tasks efficiently.

The theory expounded in this paper elucidates this rich structure and yields a proof of correctness of the algorithm. It may also be of independent interest, as a set of well-knit graph-theoretic facts.

Key words: Maximum matching problem, efficient algorithms, alternating paths, augmenting paths, blossoms, double depth first search
MSC2000 subject classification: 68Q25, 68W40
OR/MS subject classification: Primary: Analysis of Algorithms; secondary: Data structures

---

* Supported in part by NSF Grant CCF-2230414.

1. Introduction The following quote, from Lovász and Plummer's classic book [22], pg. 12, provides a nice backdrop for the work reported in this paper:

The concept of an alternating path, although quite simple, is one of the most important in all of matching theory.

For the significance of this notion in the design of efficient algorithms for matching, as well as the parallel development of the notion of an augmenting path for flow algorithms, we refer the reader to [22, 1]. The computational importance of *minimum length* augmenting paths was first recognized by Dinitz [5], in the context of flow theory, and this basic idea gave rise to several efficient maximum flow algorithms, see [1]. Independent and simultaneous works, by Hopcroft and Karp [15] and Karzanov [19], studied minimum length augmenting paths in the context of matching, and used this notion to give the most efficient algorithm of its time for maximum matching in bipartite graphs; see Section 10 for improvements obtained in recent years.

Edmonds [7] defined the notion of blossoms and used it to give the first polynomial time algorithm for finding a maximum matching in general graphs. His proof of correctness was built around graph-theoretic facts, which formalize the manner in which augmenting paths traverse blossoms and their complex nested structure.

The most efficient known algorithm for general graph matching is due to Micali and Vazirani [24], obtained in 1980[1]. It resorts to finding minimum length augmenting paths and uses the scheme proposed in [15, 19], see Section 1.2. The blossoms it finds are special and are defined in Section 8; in contrast, Edmonds' blossoms do not take into consideration length information and are therefore inadequate for the purpose of finding minimum length augmenting paths. The description of this algorithm given in [24], via a pseudo-code, is complete and error-free[2]. However, the paper did not attempt a proof of correctness.

A proof of correctness was attempted in [34]. That paper correctly recognized the fact that an elaborate, new theory of alternating paths and blossoms, from the perspective of minimum length paths, was called for to give such a proof. As detailed in Section 1.1, although that paper made some important contributions, it had serious shortcomings.

The current paper completes the task started in [34] by presenting the pertinent theory in full detail and using it to give the first complete and correct proof of the MV algorithm. In Section 1.1 we state the new ideas underlying this proof. Considering the special status of the maximum matching problem within the theory of algorithms (see below), it was not appropriate to leave its most efficient known algorithm in an essentially unproven state. Hence the investment of (substantial) effort to produce the current paper, notwithstanding the lapse of considerable time since the publication of the algorithm.

In the case of bipartite graphs, minimum length alternating paths from an unmatched vertex to a matched vertex can be of one parity only, either even or odd. Consequently, such paths possess an elementary property, called *breadth first search (BFS) honesty* in this paper: Let p be a minimum alternating path from unmatched vertex f to v and let u lie on p. Then the part of p from f to u is a minimum alternating path from f to u, and not any longer. In the presence of this property, a straightforward alternating BFS suffices for executing a phase[3] in linear time.

In general graphs, the existence of minimum length alternating paths of both parities from an unmatched vertex to a matched vertex leads to a new difficulty, whose origin lies in the fact that such paths are not BFS honest. In the situation described above, assume that p is a minimum *even*

[1] Section 10 puts this fact in context, by comparing with recent improvements in running times of other combinatorial optimization problems.

[2] For ease of comparison of the pseudo-code given in [24] to the description of the algorithm presented in the current paper, we note that a key algorithmic structure, called "petal" in the current paper, was called "blossom" in [24].

[3] See Section 1.2 for a definition.

length alternating path from f to v and u lies on p. The issue is that the part of p from f to u can be *arbitrarily longer* than a minimum path from f to u, of either parity. This happens because all minimum paths from f to u contain v at an *odd* length, see Example 2 in Section 3.2.

As a result, the following fundamental difficulty arises: For finding a minimum augmenting path in the graph, we need to find arbitrarily long paths to intermediate vertices, even though the latter do admit short paths, see Section 3. As such, this appears to call for an exponential time[4] algorithm. How then does the MV algorithm accomplish this task within the same time as bipartite graphs, i.e., linear time for a phase.

The theory expounded in this paper shows how the underlying structure offers several different "footholds" for the key tasks which the algorithm needs to perform in order to home in on a solution quickly. The bottom line is that although minimum length alternating paths are not BFS honest, they are not arbitrarily BFS dishonest. First, Theorem 2 uses the notion of tenacity to carve out an important case in which vertex u must be BFS honest on path p from f to v. Second, even if vertex u is not BFS honest on p, it is BFS honest on p w.r.t. a special vertex, called $\text{base}(u)$. In turn, even if $\text{base}(u)$ is not BFS honest on p, it is BFS honest on p w.r.t. $\text{base}(\text{base}(u)) = \text{base}^2(u)$, and so on, see Theorems 5 and 6.

Consistent with these simple-looking rules, a myriad situations can arise — illustrated via the many examples given in this paper — and the algorithm needs to efficiently handle all of them. Indeed, the structure expounded in this paper lays bare a stark contrast: *on the one hand, the extreme complexity of the problem being handled by the algorithm, and on the other, the simplicity of the algorithm itself.*

Matching has had a long and distinguished history within graph theory and combinatorics, spanning more than a century and a half [22]. Its algorithmic history is equally long, dating back to the mid-nineteenth century work of Carl Gustav Jacobi on the bipartite case, as mentioned in [37]. Its exalted status in the theory of algorithms[5] arises from the fact that its study has yielded quintessential paradigms and powerful algorithmic techniques, which form the foundation of the modern theory of algorithms, as we know it today. These include definitions of the classes $\mathcal{P}$ [8] and $\#\mathcal{P}$ [32], the primal-dual paradigm [20], the equivalence of random generation and approximate counting for self-reducible problems [17], characterizing facets of the convex hull of solutions to a combinatorial problem [7], the canonical paths argument in the Markov chain Monte Carlo method [16], and the Isolation Lemma [27, 36].

1.1. Overview and Contributions In this section, we will state the contributions of [34] and point out the nature of its shortcomings due to which the current paper was called for. We will also state the contributions of the current paper and discuss the ideas that led to the current proof.

A number of structural notions and definitions need to be given to state the algorithm and the proof. Rather than stating them all up-front, we have spread them over sections in which they are put to use for the first time. The elementary ones, pertaining to minimum length alternating paths, are given in Section 3.1. Section 4 gives definitions required for stating the algorithm, primary among them being tenacity of vertices and edges, and the classification of edges into props and bridges.

At a high level, the algorithm involves *three main ideas*:

1. **DDFS:** The new graph search procedure called *double depth first search (DDFS)* is described in Section 2 in a completely self-contained manner; it can be read without reading the rest of the paper.

[4] Recall that the problem of finding long paths, such as a Hamiltonian path, is NP-hard.

[5] See Section 10 for the role played by matching in game theory and economics.

2. **Marking petals appropriately** so paths through them can be found quickly. Petals are defined in Definition 13 and Section 5.2 provides details on marking them. Remark 1 highlights how this way of marking petals helps achieve the promised running time.

3. **Synchronization of events:** This is described in Section 9.1; in particular, see Remark 8.

Section 8 gives the central notions of base of a vertex and base of a blossom; these are essential to the proof of correctness of the algorithm and for clarifying the "footholds" mentioned in the Introduction. As described in Section 8.1, in order to define the base of a vertex v, we need to show that the set $B(v)$, defined in Definition 19, is a singleton. However, its proof requires the notion of a blossom and its associated properties. On the other hand, blossoms can be defined only after defining the base of a vertex. Therefore, we are faced with a severe "chicken-and-egg" problem.

Our proof resolves this problem by carrying out an induction on tenacity. This is done in the central structural theorem, Theorem 3. The induction basis, for the lowest tenacity vertices in the graph, is given in Section 8.2. It shows that for such vertices v, $B(v)$ is a singleton; this unique vertex is defined to be base(v). The proof of this fact is not straightforward and is accomplished by using DDFS in an appropriate setting. Indeed, one of the main innovations of the current paper is the use of this procedure, not only in the algorithm, but also for its proof.

Once the base of the lowest tenacity vertices is well-defined, the blossoms containing these vertices can be defined and properties of these blossoms and properties of paths traversing through them can be established. In the induction step, which is carried out in Section 8.3, these facts are established for higher tenacity vertices.

As mentioned above, the heart of the proof involves using DDFS in an appropriate setting; the latter is a graph obtained from the given graph. The graph in which the induction basis is carried out is far simpler than the one in which the induction step is carried out, because the latter contains blossoms defined in the previous steps of the induction. Because of this simplicity, the induction basis plays the role of a crucible in which proof techniques can be developed with relative ease for the various claims. Furthermore, as detailed below, we can determine the "correct" order in which these proofs need to be carried out. In the proof of the induction step, we first apply a transformation on previously defined blossoms, see Definition 28, after which the structure of the proofs becomes similar to those in the induction basis and these proofs are omitted, unless they contain a significant new insight.

The following steps were taken to render the algorithm easier to comprehend:

1. For ease of comprehension, DDFS has been described in the simpler setting of a directed, layered graph H. In the algorithm, DDFS is run on the original graph G. However, describing DDFS on G is too cumbersome; this was done in [34]. Instead, we provide a mapping from G to H in Section 5.3 via which the reader can easily trace the steps DDFS executes in G. Interestingly, this cosmetic-looking idea had a far-reaching conceptual consequence, as discussed below.

2. To further help the reader, in the many illustrative examples given, the distance of vertices from unmatched vertices is proportional to their minlevels, as would be the case in the corresponding layered graph H. Thus the unmatched vertices belong to the lowest level.

3. Wherever possible, procedures are described in plain English. For example, DDFS is described in English in Section 2; readers who prefer to understand it via a pseudocode can find it in [24, 34].

4. The structural definitions given in this paper are of two types: purely graph-theoretic ones and algorithmic ones; the latter depend on the manner in which ties are resolved in a run of the algorithm. Whereas the former include base and blossom, the latter include petal and bud, see Section 5.2. We have demarcated these two types of definitions and in Section 9.2 we have pointed out relationships between them.

At a high level, [34] also accurately identified the interplay between the notions of tenacity, base, blossom and bridge. However, the actual definitions given for some of these notions were incorrect. For example, in [34], the central notion of base of a vertex was defined for any vertex of finite tenacity. However, it turns out that there may be vertices of finite tenacity which have no base, e.g., see Example 14. In the current paper, base has been defined only for vertices of *eligible tenacity*, see Definition 17 for this notion. Theorem 3 in [34] "proves" that every vertex of finite tenacity has a well-defined base – it is obviously incorrect. Furthermore, Theorem 3 in [34] is incorrect even for vertices of eligible tenacity.

The errors in the proofs given in [34] can be traced back to three reasons; we provide them below together with the ways we get around them in the current paper.

1. [34] failed to identify the "chicken-and-egg" problem stated above and tried to "prove" several facts in a stand-alone manner. As mentioned above, the current paper rectifies this problem by carrying out an induction on tenacity, in Theorem 3, proving all these fact for lower tenacity vertices before moving to higher tenacity vertices.

2. The expository idea of describing DDFS in the simpler setting of a layered graph had an unexpected consequence: it led to two breakthrough ideas for the proof. First, use of the power of DDFS, and second, doing the proof in the simpler setting of graphs H'_m and H'_t. As a result of the latter, we avoid dealing with the debilitating complexity of the original graph $G=(V,E)$, unlike [34].

The induction basis is carried out in graph H'_m. Then, the non-trivial definitions of pred_t and pred^*_t, given in Definition 28, help finesse the lower tenacity blossoms, which were defined in previous induction steps, to yield the graph H'_t in which the induction step is carried out. All the facts established are then"ported" back to G using the mapping between these graphs and G, see Sections 8.2.1 and 8.3.1.

These ideas qualitatively changed[6] the nature of the proofs: instead of dealing with individual alternating paths, which can get arbitrarily complicated, leading to incorrect proofs, we now deal with the very precise and potent information provided by the DDFS Certificate, as proven in Theorem 1.

3. The mindset in [34] was the following: The culminating fact which needed to be proven was the existence of a bridge of the right tenacity on a maxlevel path[7], and a variety of other structural facts needed to be established prior to that. This order turns out to be incorrect. As mentioned above, in the current paper, the order is dictated by the proof of the induction basis, carried out in graph H'_m; its simple structure makes transparent the "right" order of implications, see Section 8.2. This is the right order for the induction step as well, see Section 8.3. In particular, the existence of a bridge is the first, and not the last, fact we establish in the induction basis, as well as in the induction step.

1.2. Running Time and Related Papers The simplest scheme for finding a maximum matching is to start with the empty matching, and iteratively find augmenting paths and augment the current matching. When there is no such path, the matching must be maximum; this is shown in [2]. In order to improve the running time, [15, 19] proposed finding multiple augmenting paths in each iteration as stated in Definition 1.

Definition 1. (*Phase*) In a graph $G=(V,E)$ with matching M, a phase consists of finding a maximal set of disjoint minimum length augmenting paths and augmenting M along all paths found.

[6] In a sense, the current proof owes its existence to a chance event: our attempt at simplifying the exposition of DDFS.

[7] This viewpoint is not unreasonable, e.g., see Section 9.4.

As shown in [15, 19], $O(\sqrt{n})$ phases suffice[8] for finding a maximum matching. These papers also show how to implement a phase in $O(m)$ time in a bipartite graph, thereby getting a total running time of $O(m\sqrt{n})$. The MV algorithm executes a phase in linear time on the RAM model. Its total running time is $O(m\sqrt{n})$ on the RAM model and $O(m\sqrt{n}\cdot\alpha(m,n))$ on the pointer model (see Theorem 8 for details).

We note that small theoretical improvements to the running time, for the case of very dense graphs, have been given in recent years: $O(m\sqrt{n}\log(n^2/m)/\log n)$ [14] and $O(n^w)$ [26] where w is the best exponent of n for multiplication of two $n\times n$ matrices. The former improves on MV for $m=n^{2-o(1)}$ and the latter for $m=\omega(n^{1.85})$. However, the latter algorithm involves a large multiplicative constant in its running time which comes from the use of fast matrix multiplication as a subroutine, thereby making the small improvement in the exponent not very meaningful.

Prior to [24], [9] had used the idea of finding augmenting paths in phases to obtain an $O(n^{2.5})$ maximum matching algorithm. However, their algorithm is extremely complicated and its correctness is hard to ascertain, in particular because there is no journal version of that result.

Subsequent to [24], [11] gives an efficient scaling algorithm for finding a minimum weight matching in a general graph with integral edge weights and at the end of the paper, it claims that the unit weight version of their algorithm achieves the same running time as MV; see also [12]. The rest of the history of matching algorithms is very well documented and will not be repeated here, e.g., see [22, 34].

2. Double Depth First Search (DDFS) This section is fully self-contained and describes the procedure of double depth first search (DDFS). For ease of comprehension, we have presented DDFS in the simplified setting of a directed, layered graph H. In the MV algorithm, DDFS is run on the original graph G, which is far more complex. In G, DDFS terminates with either a new blossom — more precisely, a new petal — or the existence of a new augmenting path. These correspond to Case 1 and Case 2, detailed below. We will provide an explicit mapping between the two settings in Section 5.3.

The input to DDFS is a *directed, layered graph* $H=(V,E)$. V is partitioned into $h+1$ layers, for some $h>0$. The layers are numbered from 0 to h and are named $l_0,\ldots,l_h$, with l_0 being the *lowest layer* and l_h the *highest layer*. The layer number of a vertex v is denoted by $l(v)$. We will assume that for each $v\in V$, $l(v)$ is easily available; in fact, it can be obtained in unit time. If $l(v)<l(u)$, then we will say that v *is deeper than* u. Each directed edge $(u,v)\in E$ runs from a higher to a lower layer, not necessarily consecutive, i.e., $l(u)>l(v)$. V contains two special vertices, r and g, for *red* and *green*, not necessarily in the same layer, and neither of them in l_0. See Figure 1 for a layered graph with $h=7$.

At a high level, the objective of DDFS is to grow two DFS trees, T_r and T_g, rooted at r and g, respectively, in such a way that T_r and T_g *share at most one vertex* and the deepest vertex (vertices) in the two trees is (are) "as deep as possible". Furthermore, this needs to be done in time that is linear in the sum of the sizes of the two trees. Because of the DDFS Requirement, stated below, it is trivial to grow any one tree very deep, all the way to the lowest layer, l_0. However, this will not achieve the more interesting and useful objective stated above. For that, we grow each tree in such a way that it does not "block off" the other tree, by growing them in a highly coordinated manner; the latter is the main point of DDFS.

We require that H satisfies:

DDFS Requirement: Starting from every vertex $v\in V$, there is a path to a vertex in layer l_0.

[8] As is standard, n denotes the number of vertices and m the number of edges in the given graph.

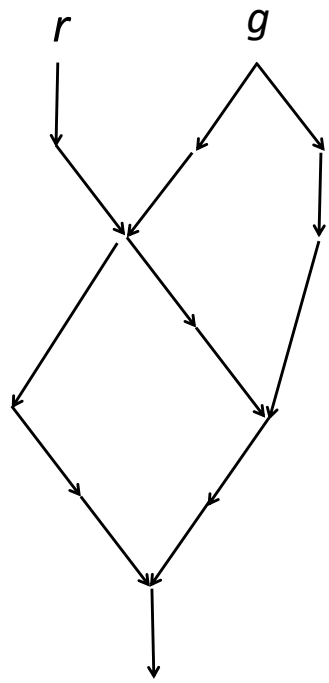

FIGURE 1. Layered graph H with r and g in layer l_7.

Vertex v will be called a *bottleneck* if every path from r to l_0 and every path from g to l_0 contains v; v is allowed to be r or g or a vertex in layer l_0. Let p be a path from r or g to layer l_0. Since layer numbers on p are monotonically decreasing, if there is a bottleneck, the one having highest level must be unique. It will be called the *highest bottleneck* and we will denote it by b. If H has a bottleneck, we will say that we are in *Case 1*. Otherwise, there must be distinct vertices r_0 and g_0 in layer l_0 such that there are vertex-disjoint paths from r to r_0 and g to g_0; this will be called *Case 2*.

As stated above, in the MV algorithm, these two cases correspond to the creation of a new petal and the discovery of a new augmenting path, respectively. In the graph of Figure 1, DDFS will terminate in Case 1, with bottleneck b, as shown in Figure 2. In the graph of Figure 3, which differs from the graph of Figure 1 only in the two edges going from c to g_0, DDFS terminates in Case 2, with disjoint paths from r to r_0 and g to g_0.

In Case 1, let V_b (E_b) be the set of all vertices (edges) that lie on all paths from r or g to b. In Case 2, let E_p be the set of all edges that lie on all paths starting from r or g and ending at r_0 or g_0.

The objective of DDFS: The first objective of DDFS is to determine which of these two cases holds. Additional objectives of DDFS in the two cases are:

Case 1: DDFS needs to find the highest bottleneck, b, and partition the vertices in $V_b - \{b\}$ into two sets S_R and S_G, called the *red set* and *green set* respectively, with $r \in S_R$ and $g \in S_G$. These sets should satisfy:

1. There is a path from r to b in $S_R \cup \{b\}$ and a path from g to b in $S_G \cup \{b\}$.
2. There are two spanning trees, T_r and T_g, in $S_R \cup \{b\}$ and $S_G \cup \{b\}$, and rooted at r and g, respectively. Furthermore, DDFS needs to find such a pair of trees.

Case 2: DDFS needs to find distinct vertices r_0 and g_0 in layer l_0, and vertex disjoint paths from r to r_0 and g to g_0. In this case, as soon as DDFS finds these two paths, it halts, even if it has not traversed all edges of E_p, since a new augmenting path has already been found. Let $E'_p \subseteq E_p$ denote the edges which DDFS has actually traversed.

The two DFS trees: DDFS involves the coordinated growth of two DFS trees, the *red tree* T_r and the *green tree* T_g, rooted at r and g, respectively. At each point in the algorithm, each tree has a well-defined *center of activity*, i.e., the vertex it is currently exploring. These are denoted by C_r and C_g and are initialized to the two roots, r and g, respectively. When a center of activity is at a vertex u and is ready to move, it must be the case that the color of u is the same as that of the center of activity. If the center moves to a vertex v, the edge (u,v) is assigned to the corresponding tree and given the color of the center.

Therefore, each edge (u,v) has the same color as that of u, i.e., all edges out of u will get the color of u. Note however that the color of v may be different from that of u. Figure 2 shows T_r and T_g after DDFS has been performed on the graph of Figure 1. T_r consists of broken edges and T_g consists of solid edges; the edge from b to l_0 is in neither tree. Note that b is in both trees and gets neither color.

At termination, DDFS provides the following certificate.

Definition 2. (DDFS Certificate) In Case 1, for every vertex $v \in V_b - \{b\}$, if v is red, there is a path from r to v in T_r and a disjoint path from g to b in T_g. And if v is green, there exists a path from g to v in T_g and a disjoint path from r to b in T_r. In Case 2, there are vertex disjoint paths from r to r_0 and g to g_0 having colors red and green, respectively.

Running time: The running time of DDFS needs to be $O(|E_b|)$ in Case 1 and $O(|E'_p|)$ in Case 2.

Coordinated growth of the two trees: We will first describe aspects of DDFS in which the two trees function as "normal" DFS trees in a directed graph, and then we will describe their coordination; in particular, the coordination determines, at each step, which tree grows. Initially, all vertices, other than r and g, are marked "unvisited" and all edges are marked "unexplored".

Every vertex in $T_r \cup T_g$, other than r, g, and b, has a unique *parent*; r and g have no parent and b has two parents, one of each color. Assume that $C_r = u$ and it is T_r's turn to grow. If so, T_r picks an unexplored edge, say (u,v), out of u. If v is already marked "visited" and $C_g \neq v$, then T_r picks another unexplored edge out of u. If v is marked "unvisited", then v is marked "visited", u is designated the parent of v, and C_r moves to v. The last case, i.e., v is already marked "visited" and $C_g = v$ is dealt with below. When all outgoing edges from u have been explored, C_r backtracks from u to its parent if $u \neq r$; the case $u = r$ is dealt with below. The growth of T_g is analogous. Since H is acyclic, the trees have no back edges.

We next describe the coordination between the two trees. We will adopt the (arbitrary) convention that C_r will "try to keep ahead of" C_g, and C_g will "try to catch up". Following this convention, the moves of C_r are as follows: If $l(C_r) > l(C_g)$, then C_r keeps moving until the first time that $l(C_r) \leq l(C_g)$. If $l(C_r) = l(C_g)$ and $C_r \neq C_g$, then C_r moves one step and stops; at this point, $l(C_r) < l(C_g)$. If $C_r = C_g$, the two centers of activity have met and this case is described below. The moves of C_g are as follows: If $l(C_r) < l(C_g)$, then C_g keeps moving until the first time that $l(C_g) \leq l(C_r)$ and then it stops.

When the two centers of activity meet: As stated above, when one of the centers of activity traverses an edge, the edge is assigned to the corresponding tree and is assigned its color.

However, the assigning of color to a vertex is not so straightforward and is not done in a greedy manner. Indeed, a vertex v may first be added to one tree and later this decision may be reverted;

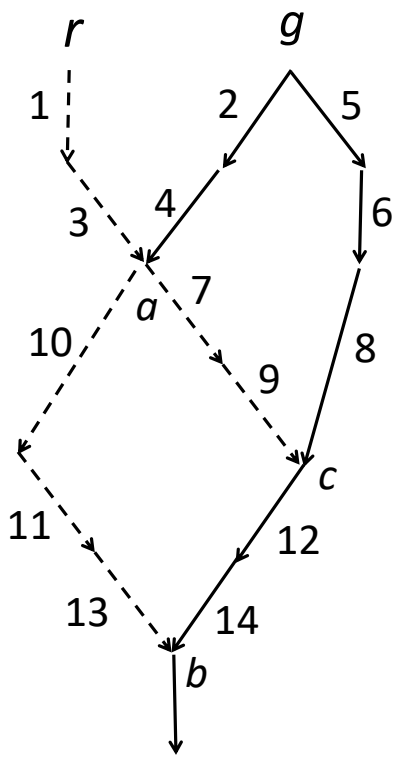

FIGURE 2. DDFS executed from r and g terminates in Case 1 with bottleneck b. Edge numbers indicate the order of traversal of the edges.

this happens if C_r and C_g meet at v[9]. We will adopt the convention that when this happens, first C_g backtracks and tries to find an alternative path that is as deep as v. If it fails, then C_g occupies v and C_r tries to find an alternative path that is as deep as v. If C_r also, fails, then it must have backtracked all the way to the root r and DDFS terminates.

We will explain these moves in detail via Figure 2. In this figure, the numbers on the edges indicate the order in which they are added to the two trees. Observe that the two centers of activity meet for the first time at a. At this point, DDFS needs to determine if a is the highest bottleneck, and if not, then which of the trees can find an alternative path at least as deep as a, so search may resume. By the convention established above, C_g tries first. After it backtracks all the way to g, it traverses edges number 5 and 6 and arrives at a vertex that is as deep as a, and DDFS resumes.

The two centers of activity meet for the second time at vertex c. This time, C_g backtracks all the way to g, without finding an alternative path. As per our convention, C_g now occupies c and C_r tries to find an alternative path as deep as c.

However, at this stage, we need to introduce an important notion, namely the pointer *Barrier*. Its purpose is to prevent C_g from backtracking from a vertex more than once. At the start of DDFS, the Barrier is initialized to g. At this stage, since C_g has backtracked from c all the way to g, i.e., the current position of the Barrier, it is now moved to c.

[9] Observe that either of the trees could have arrived at v first. This happens despite our convention that C_r keeps ahead of C_g — the reason is that C_g may have used a long edge to arrive at v before C_r. In Figure 2, this happens when the two trees meet at vertex c.

Next, the two centers of activity meet at b. By our convention, b is first given to T_r and C_g attempts to find an alternative path. However, it backtracks all the way to the Barrier, which is currently at c, without success. At this point, the Barrier is moved to b, b is given to T_g and C_r attempts to find an alternative path. However, it backtracks all the way to r without finding an alternative path. At this point, we conclude that b is the bottleneck. In general, when C_r backtracks all the way to r, DDFS terminates in Case 1 and the current meeting point is declared the bottleneck.

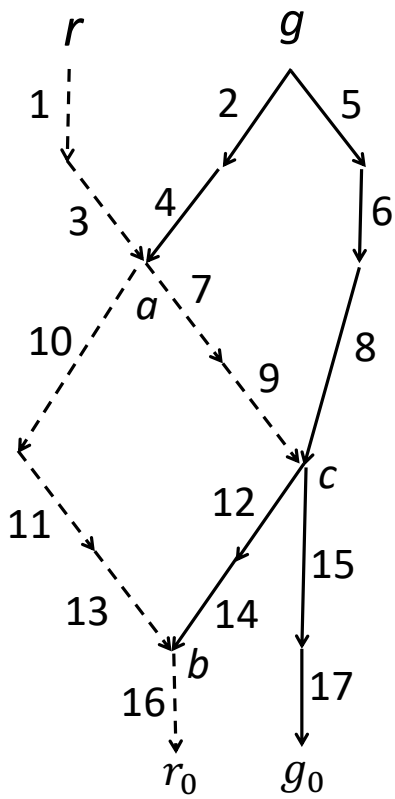

FIGURE 3. DDFS executed from r and g terminates in Case 2.

In Figure 3, after backtracking from b, C_g does manage to find an alternative path as deep as b, when it explores edge number 15. At this point, DDFS resumes and C_r reaches r_0 in layer l_0 and C_g reaches g_0 in that layer, hence terminating in Case 2.

THEOREM 1. *DDFS accomplishes the objectives stated above in the required time.*

Proof : In Case 1, tree T_r contains paths consisting of red colored vertices from r to b and from r to each red vertex. A similar claim holds about tree T_g. In Case 2, there is a path consisting of red colored vertices from r to r_0 in tree T_r and there is a path consisting of green colored vertices from g to g_0 in tree T_g. Therefore, the DDFS Certificate holds.

Finally, it is easy to see that each edge of H is explored by at most one tree and if so, only once. Clearly, T_r backtracks from each vertex at most once and because of the Barrier, the same holds for T_g as well. The theorem follows. □

3. Elementary Definitions and a Fundamental Notion In Section 3.1, we will present some elementary definitions pertaining to minimum length augmenting paths. Using these definitions, in Section 3.2 we present the property of breadth first search honesty, due to which an alternating BFS works in bipartite graphs, yielding a linear time algorithm for a phase. The lack of this property in non-bipartite graphs necessitates a much more complex algorithm.

3.1. Elementary Definitions A matching M in an undirected graph $G=(V,E)$ is a set of edges no two of which meet at a vertex. Our problem is to find a matching of maximum cardinality in the given graph. *Henceforth all definitions will be w.r.t. a fixed matching M in G.* Edges in M will be said to be *matched* and those in $E-M$ will be said to be *unmatched*. Vertex v will be said to be matched if there is a matched edge incident at it and unmatched otherwise.

An *alternating path* is a simple path whose edges alternate between M and $E-M$, i.e., matched and unmatched. An alternating path that starts and ends at unmatched vertices is called an *augmenting path*. Clearly the number of unmatched edges on such a path exceeds the number of matched edges on it by one[10]. Its significance lies in that flipping matched and unmatched edges on such a path leads to a valid matching of one higher cardinality. Edmonds' matching algorithm operates by iteratively finding an augmenting path w.r.t. the current matching, which initially is assumed to be empty, and augmenting the matching. When there are no more augmenting paths w.r.t. the current matching, it can be shown to be maximum.

The MV algorithm finds augmenting paths in phases as proposed in [15, 19]. In each *phase*, it finds a maximal set of disjoint minimum length augmenting paths w.r.t. the current matching and it augments along all paths. [15, 19] show that only $O(\sqrt{n})$ such phases suffice for finding a maximum matching in general graphs. The remaining task is designing an efficient algorithm for a phase.

Definition 3. (Length of Minimum Length Augmenting Path) Throughout, l_m will denote the length of a minimum length augmenting path in G; if G has no augmenting paths, we will assume that $l_m=\infty$.

Definition 4. (*Evenlevel and oddlevel of vertices*) The evenlevel (oddlevel) of a vertex v, denoted by evenlevel(v) (oddlevel(v)), is defined to be the length of a minimum even (odd) length alternating path from an unmatched vertex to v; moreover, each such path will be called an evenlevel(v) (oddlevel(v)) path. If there is no such path, evenlevel(v) (oddlevel(v)) is defined to be ∞.

We will typically denote an unmatched vertex by f. Its evenlevel is zero and its oddlevel is the length of the shortest augmenting path starting at f; if no augmenting path starts at f, oddlevel(f) $=\infty$. The length of a minimum length augmenting path w.r.t. M is the smallest oddlevel of an unmatched vertex.

Definition 5. (*Maxlevel and minlevel of vertices*) For a vertex v such that at least one of evenlevel(v) and oddlevel(v) is finite, maxlevel(v) (minlevel(v)) is defined to be the bigger (smaller) of the two.

Definition 6. (*Outer and inner vertices*) A vertex v with finite minlevel is said to be *outer* if evenlevel(v) $<$ oddlevel(v) and *inner* otherwise.

Definition 7. (Odd and even w.r.t. p) Let p be an alternating path from unmatched vertex f to v and let u lie on p. The *length* of path p, denoted by $|p|$, is the number of edges on p. The part of p from f to u will be denoted by $p[f \text{ to } u]$ and $p[f \text{ to } u)$ will denote the part of p from f to the vertex just before u. Other combinations are self-explanatory. We will say that vertex u is *even w.r.t. p* if $|p[f \text{ to } u]|$ is even and it is *odd w.r.t. p* if $|p[f \text{ to } u]|$ is odd.

[10] Observe that if $M=\emptyset$, then any edge is an augmenting path, of length one.

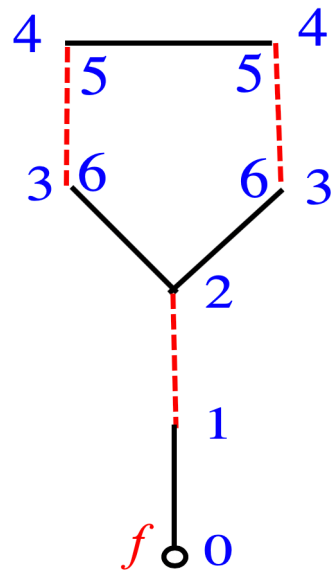

FIGURE 4. The evenlevels and oddlevels of vertices are indicated; missing levels are ∞.

EXAMPLE 1. In the figures hereafter, matched and unmatched edges are drawn with broken and solid lines, respectively. Additionally, unmatched vertices are drawn with a small circle, e.g., vertex f in Figure 4. The numbers in this figure indicate the evenlevels and oddlevels of vertices, with missing numbers being infinity.

3.2. The Notion of BFS Honesty Let p be an alternating path from unmatched vertex f to v. We will say that p is a *minimum alternating path* if $|p| = \text{evenlevel}(v)$ ($|p| = \text{oddlevel}(v)$) if $|p|$ is even (odd).

Breadth first search (BFS) honesty is the following property: Let p be a minimum alternating path from unmatched vertex f to v and let u lie on p. Then $p[f \text{ to } u]$ is a minimum alternating path from f to u. Bipartite graphs satisfy this property, and as a consequence, a straightforward alternating BFS suffices for finding minimum augmenting paths, see Section 4, or for a complete description, see Section 2.1 in [34].

Surprisingly enough, this elementary property does not hold in non-bipartite graphs, as illustrated in Example 2. This basic difference arises because in bipartite graphs, minimum length alternating paths from an unmatched vertex f to a vertex v can be of one parity only, either even or odd, but in non-bipartite graphs, they can be of both parities. As a result, the following may happen: Let p be an evenlevel(v) path from f to v and let u lie on it with $|p[f \text{ to } u]|$ being odd, say. Then $p[f \text{ to } u]$ can be arbitrarily longer than an oddlevel(u) path. The reason is that every oddlevel(u) path, say q, contains v at an odd length, and extending q to v to get an even length path will result in a self-intersecting path, see Example 2. Consequently, for the graph in Figure 5, we would need to find longer and longer odd-length alternating paths from f to w in order to find minimum length alternating paths from f to other vertices, e.g., u and v.

In summary, the following fundamental difficulty arises: For finding a minimum length augmenting path, we need to find arbitrarily long paths to intermediate vertices, even though the latter do admit short paths. As such, this appears to call for an exponential time algorithm. Recall that finding short paths is easy and long paths is hard, e.g., Hamiltonian path is NP-hard.

EXAMPLE 2. In Figure 5, oddlevel(w) = 7. An evenlevel(v) path is shown in this figure. Observe that w occurs at a length of 11 on this path. Also observe that v occurs at an odd length on the oddlevel(w) path. It will be instructive for the reader to find an evenlevel(u) path; observe that w occurs at a length of 9 on it.

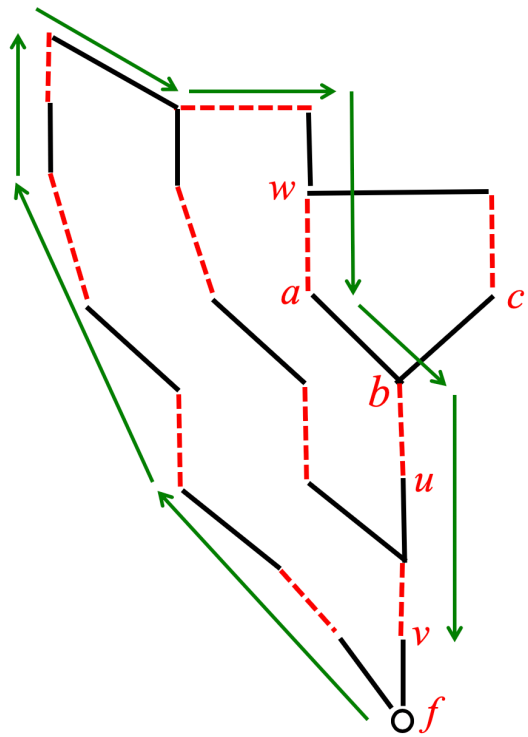

FIGURE 5. Vertices w, a, b and u are not BFS honest on the evenlevel(v) path shown via arrows.

4. Some Essential Definitions As mentioned in the Introduction, in order to implement a phase in linear time in non-bipartite graphs, we need to exploit the elaborate structure offered by minimum length alternating paths. In this section, we present some facts that are absolutely necessary to describe the MV algorithm. The proof of correctness of the algorithm requires additional structural properties, presented in Section 8.

DEFINITION 8. (*Tenacity of vertices and edges*) Define the tenacity of vertex v, tenacity$(v) =$ evenlevel$(v) +$ oddlevel(v). If (u, v) is an unmatched edge, then tenacity$(u, v) =$ evenlevel$(u) +$ evenlevel$(v) + 1$, and if it is matched, tenacity$(u, v) =$ oddlevel$(u) +$ oddlevel$(v) + 1$.

The notion of tenacity is central to the structural facts that follow. Clearly tenacity$(f) \geq l_m$ for an unmatched vertex f, see Definition 3 for the notion of l_m. Furthermore, tenacity$(f) = l_m$ if and only if f participates in a minimum length augmenting path.

DEFINITION 9. (Minimum tenacity of a vertex in G) Throughout, t_m will denote the tenacity of a minimum tenacity vertex in G.

Clearly $t_m \leq l_m$. If $t_m = l_m$, the situation is particularly simple, since there are no blossoms and essentially the bipartite graph algorithm works for executing a phase. Henceforth we will assume that $t_m < l_m$.

EXAMPLE 3. In Figure 6, the tenacities of vertices are marked. They are $\alpha = 13$, $\beta = 15$, $\gamma = 17$ and $\delta = \infty$. In Figure 7, the tenacities of edges are marked. They are $\alpha = 13$, $\beta = 15$ and $\gamma = 17$.

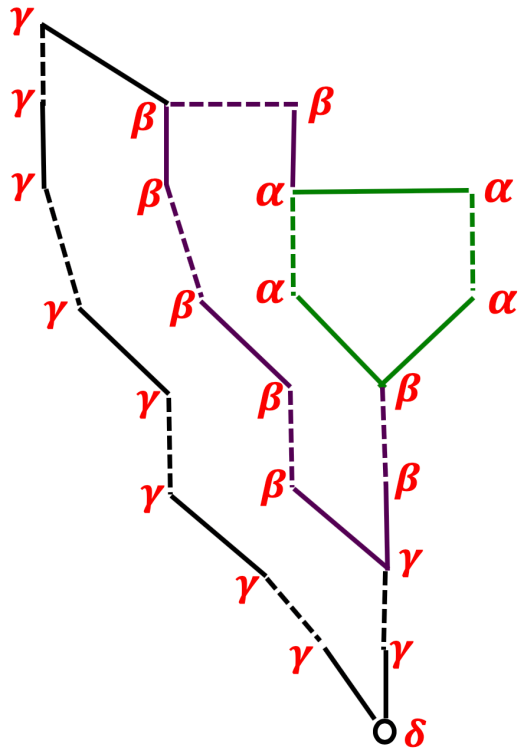

FIGURE 6. The tenacity of vertices is indicated; here $\alpha = 13$, $\beta = 15$, $\gamma = 17$ and $\delta = \infty$.

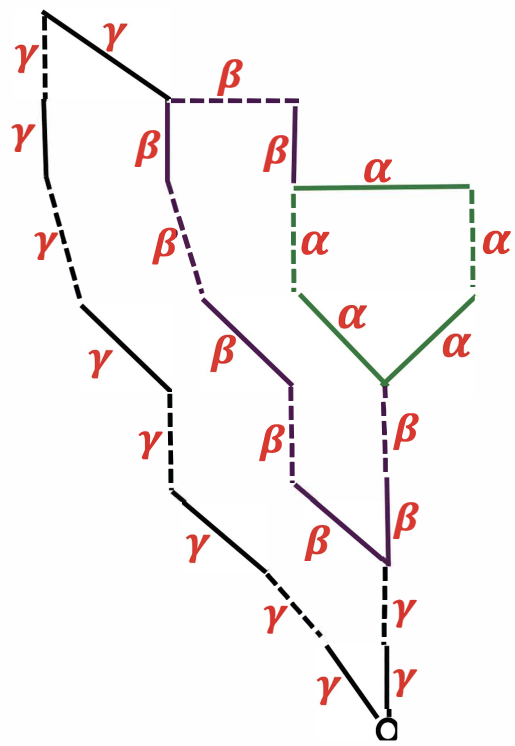

FIGURE 7. The tenacity of each edge is indicated; here $\alpha = 13$, $\beta = 15$ and $\gamma = 17$.

DEFINITION 10. (*Predecessor, prop and bridge*) Consider a minlevel(v) path and let (u,v) be the last edge on it; clearly, (u,v) is matched if v is outer and unmatched otherwise. In either case, we will say that u is a predecessor of v and that edge (u,v) is a prop. An edge that is not a prop will be defined to be a bridge.

DEFINITION 11. (The relations pred and pred*) Let v be a vertex such that minlevel(v) is finite. If v is an outer vertex, it will have a unique predecessor, namely its matched neighbor; otherwise, it will have one or more predecessors. The relation pred is defined as follows: we will say that u *is* pred v if u is a predecessor of v; we will also write is as $u = \text{pred } v$. The relation pred* is the reflexive, transitive closure of the relation pred. If u is pred* v, we will also write is as $u = \text{pred}^*\ v$.

EXAMPLE 4. In Figure 5, the two horizontal edges and the oblique unmatched edge at the top are bridges and in Figure 9, (w,w') and (u,v) are bridges; the rest of the edges in these two graphs are props. In Figure 8, edge (u,v) is a bridge. This bridge is unusual because u is pred$^*v$, even though u is not pred v.

DEFINITION 12. (*The support of a bridge*) Let (u,v) be a bridge of tenacity $t \le l_m$. Then, its support is defined to be

$$\text{support}(u,v) = \{w \mid \text{tenacity}(w) = t \text{ and } \exists \text{ a maxlevel}(w) \text{ path containing } (u,v)\}.$$

EXAMPLE 5. In the graph of Figure 6, the supports of the bridges of tenacity α, β and γ are the sets of vertices of tenacity α, β and γ, respectively. In the graph of Figure 8, the tenacity of bridge[11] (u,v) is 13 and its support consists of two vertices, v and its matched neighbor. In Figure

[11] This bridge is very unusual: on the one hand neither endpoint of a bridge is a predecessors of the other and on the other hand, in the case of this bridge, one of its endpoints $u = \text{pred}^*v$, as per Definition 11.

9, the supports of the bridges (w, w') and (u, v) are all vertices of tenacity 11 and 13, respectively. Observe that in Figure 9, f is not in the support of any bridge and tenacity$(f) = \infty$.

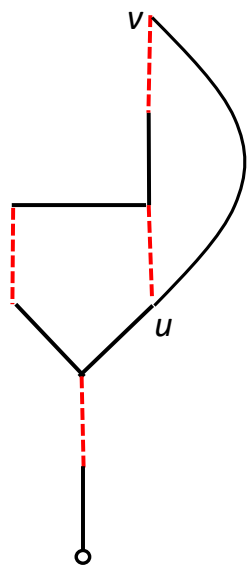

FIGURE 8. Edge (u, v) is a bridge.

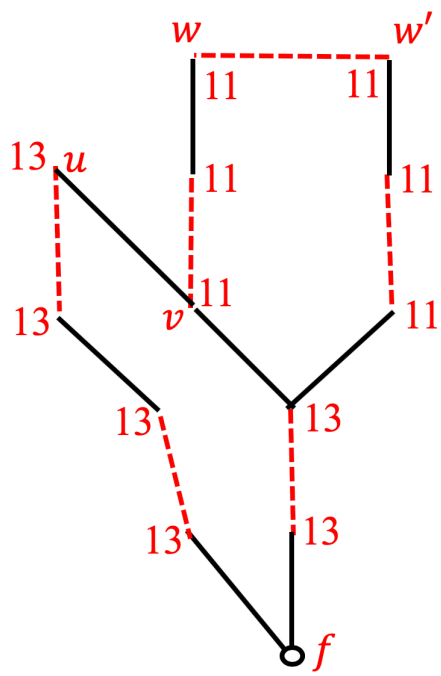

FIGURE 9. Edges (w, w') and (u, v) are bridges, with the latter being an anomaly bridge.

5. A Description of the MV Algorithm The MV algorithm executes phases as defined in Section 1.2. Each phase starts with the matching, say M, computed in the last phase. A basic task accomplished in a phase is finding the minlevels and maxlevels of all vertices of tenacity $\leq l_m$. For this purpose, the algorithm calls the procedures MIN and MAX iteratively as described below.

In this section, we have described the MV algorithm using only the definitions stated in Section 4. However, in some places, more clarity results from using notions that are defined later in the paper; if so, we have referred to the appropriate definitions. The reader is advised to get a broad idea of the algorithm on first reading and occasionally return to this section while reading the rest of the paper.

5.1. Procedures MIN and MAX At the beginning of a phase, all unmatched vertices are assigned a minlevel of 0, the rest are assigned a temporary minlevel of ∞. No vertices are assigned maxlevels at this stage. The algorithm for a phase is organized in search levels, denoted by i, starting at 0. At each search level, MIN executes one step of alternating BFS and is followed by MAX, which executes DDFS, if needed. See Algorithm 1 for a summary of the main steps.

If i is even (odd), MIN searches from all vertices, u, having an evenlevel (oddlevel) of i along incident unmatched (matched) edges, say (u, v). If edge (u, v) has not been scanned before, MIN will determine if it is a prop or a bridge as follows. If v has already been assigned a minlevel of at most i, then (u, v) is a bridge. Otherwise, v is assigned a minlevel of $i + 1$, u is declared a predecessor of v and edge (u, v) is declared a prop. Note that if i is odd, v will have only one predecessor – its matched neighbor, and if i is even, v will have one or more predecessors.

Once an edge is identified as a bridge, if MIN is able to ascertain its tenacity, say t, then the edge is inserted in the *list of bridges of tenacity* t, $Br(t)$. MIN is able to ascertain the tenacity of a bridge as long as it is not an *anomalous bridge*, as defined below; in the latter case, MAX finds the tenacity of this bridge. Task 2 in Theorem 7 proves that by the end of execution of procedure MIN at search level i, the algorithm would have identified every bridge of tenacity $2i + 1$.

Algorithm 1 **At search level i:**

1. **MIN:**

For each level i vertex, u, search along appropriate parity edges incident at u.
  For each such edge (u,v), if (u,v) has not been scanned before **then**
    If $\text{minlevel}(v) \geq i+1$ **then**
      $\text{minlevel}(v) \leftarrow i+1$
      Insert u in the list of predecessors of v.
      Declare edge (u,v) a prop.
    Else declare (u,v) a bridge,
        and if $\text{tenacity}(u,v)$ is known, insert (u,v) in $Br(\text{tenacity}(u,v))$.
  End
End

2. **MAX:**

For each edge in $Br(2i+1)$:
  Find its support using DDFS.
  For each vertex v in the support:
    $\text{maxlevel}(v) \leftarrow 2i+1-\text{minlevel}(v)$
    If v is an inner vertex, **then**
      For each edge e incident at v which is not prop,
        if its tenacity is known, insert e in $Br(\text{tenacity}(e))$.
      End
  End
End

After MIN is done, procedure MAX calls DDFS with each bridge of tenacity $2i+1$ and finds the support of this bridge. In the process, DDFS finds all vertices, v, having $\text{tenacity}(v) = 2i+1$. Since their minlevels are at most i, they are already known, and hence $\text{maxlevel}(v) = 2i+1-\text{minlevel}(v)$ can be easily computed. Clearly, if $\text{minlevel}(v)$ is $\text{evenlevel}(v)$ then $\text{maxlevel}(v)$ will be $\text{oddlevel}(v)$ and if $\text{minlevel}(v)$ is $\text{oddlevel}(v)$ then $\text{maxlevel}(v)$ will be $\text{evenlevel}(v)$.

Example 6. For each bridge in the first five figures, its tenacity gets ascertained by MIN (including the bridge (u,v) in Figure 8). We next explain the notion of an *anomalous bridge* via the graph in Figure 9. At search level 4, MIN searches from vertex u along edge (u,v) and realizes that v already has a minlevel of 3 assigned to it. Moreover, u got its minlevel from its matched neighbor. Therefore, MIN correctly identifies edge (u,v) to be a bridge. However, it is not able to ascertain $\text{tenacity}(u,v)$ since $\text{evenlevel}(v)$ is not known at this time. At search level 5, after conducting DDFS on bridge (w,w') (of tenacity 11), MAX will assign $\text{maxlevel}(v) = 8$, which is also $\text{evenlevel}(v)$. Therefore, at that time, $\text{tenacity}(u,v)$ will be ascertained to be 13 by MAX and edge (u,v) is inserted in $Br(13)$. Thus (u,v) is an anomalous bridge.

Let us explain this notion in more general terms. Let (u,v) be an unmatched bridge such that the evenlevel of one of the endpoints, say v, has not been determined at the point when MIN realizes that (u,v) is a bridge; if so v must be an inner vertex. The evenlevel of v will be determined by MAX at search level $(\text{tenacity}(v)-1)/2$ and at this point, $\text{tenacity}(u,v)$ is ascertained and the edge is inserted in $Br(\text{tenacity}(u,v))$. An important point to note in Figure 9, is that $\text{tenacity}(v) < \text{tenacity}(u,v)$. This ensures that $\text{maxlevel}(v)$ is known at search level $(\text{tenacity}(v)-1)/2$, i.e., before the search level at which bridge (u,v) needs to be processed by MAX, namely search level $(\text{tenacity}(u,v)-1)/2$.

Assume that DDFS is processing a bridge of tenacity $2i+1$ and vertex v is in its support. If v is inner and has an incident unmatched edge (u,v) which is not a prop, then it must be an anomalous bridge. MAX will ascertain its tenacity and insert it in $Br(\text{tenacity}(u,v))$. Note that $\text{tenacity}(u,v) > 2i+1$ and bridge (u,v) will need to be processed in a higher search level.

Let l_m be the length of a minimum length augmenting path in a phase. Then during search level j_m, where $l_m = 2j_m + 1$, a maximal set of such paths is found. This is described in Section 5.4.

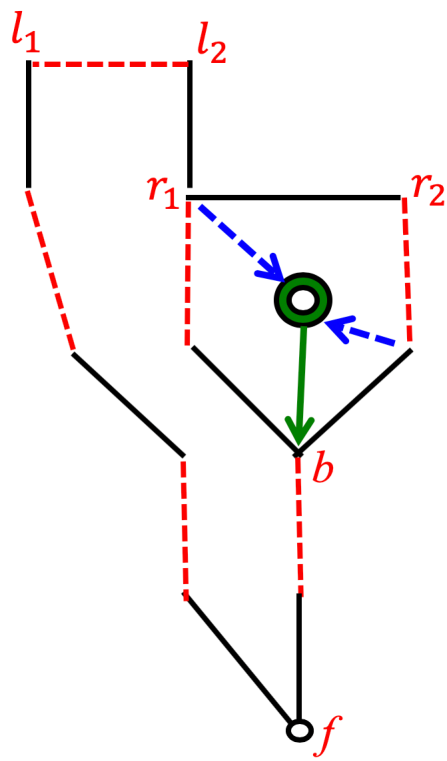

FIGURE 10. A new petal-node is created after DDFS is performed on bridge (r_1, r_2).

5.2. Petals and the Way they are Marked The notion of a petal is central to the MV algorithm; petals are found by DDFS.

DEFINITION 13. (Petal and bud) Assume that DDFS is called with a bridge (u,v) of eligible tenacity t and it terminates in Case 1. Then the bottleneck b found is called a *bud*. The set of vertices of tenacity t encountered during the current DDFS, which must lie in the support of (u,v), form a new *petal*. Formally, the petal consists of all vertices in the support of (u,v) minus the supports of all bridges (of tenacity t) processed thus far in the current search level.

Note that b must be an outer vertex and it is not included in the current petal. Observe that each vertex of eligible tenacity is included in exactly one petal and the formation of each petal is triggered by a distinct bridge[12].

EXAMPLE 7. In the graph of Figure 10, MAX will call DDFS with the bridge (r_1, r_2), which is of tenacity 9, at search level 4. The two DFSs will be rooted at r_1 and r_2, and DDFS will terminate in Case 1 with b as the highest bottleneck. The four vertices which constitute the support of bridge (r_1, r_2) form the new petal and b is the bud of this petal. Observe that b does not belong to this petal.

[12] Not every bridge will lead to the formation of a petal, e.g., see Example 10.

Marking a petal: When a new petal is found, the algorithm executes the following steps: It creates a new node, called *petal-node*; this has the shape of a doughnut in Figure 10. All vertices of the new petal point[13] to the petal-node; b is not in the petal and does not point to the petal-node. The new petal-node points to the two endpoints of its bridge, r_1 and r_2, and to its bud, b. These pointers will enable the algorithm to:

1. Skip over this petal in future DDFSs.
2. Efficiently find an alternating path through the petal, see Remark 1.

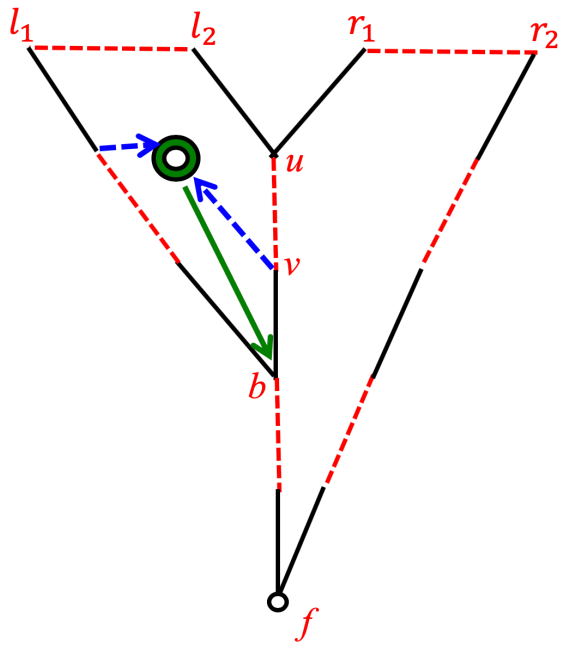

FIGURE 11. The bud formed when DDFS is performed on bridge (l_1, l_2); its bud is b.

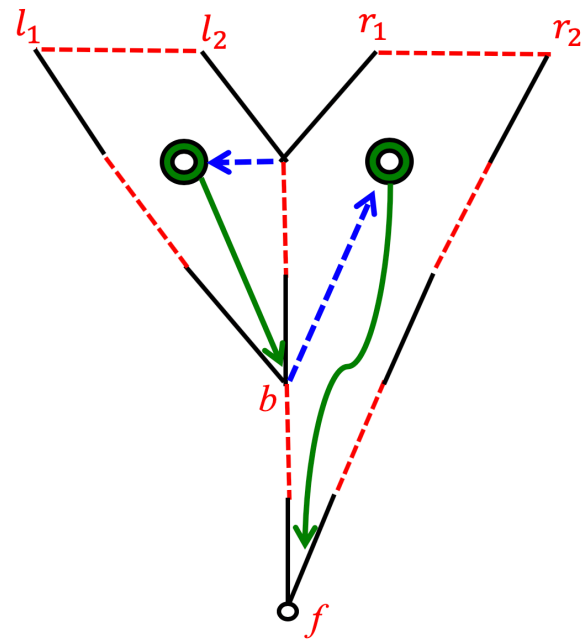

FIGURE 12. The bud and blossom formed when DDFS is performed on bridge (r_1, r_2); their bud and base is f.

DEFINITION 14. (The functions $\text{bud}(v)$ and $\text{bud}^*(v)$) Define a function $\text{bud}: V \to V$ as follows. If vertex v is in a petal then $\text{bud}(v) = b$, where the bud of this petal is b, and if v is not in a petal, then $\text{bud}(v) = v$. At any point in the execution of the algorithm, the function $\text{bud}^*(v)$ is defined recursively as follows: If $\text{bud}(v) = v$ then $\text{bud}^*(v) = v$, else $\text{bud}^*(v) = \text{bud}^*(\text{bud}(v))$. Clearly, $\text{bud}^*(v)$ will keep changing as the algorithm proceeds.

The notions of petal and bud are intimately related to the notions of blossom and base. Whereas the first pair is algorithmic — the exact petals and buds found depend on the manner in which the algorithm resolves ties — the second pair is purely graph-theoretic. The relationship between these notions is established in Section 9.2. Here we simply note that a blossom is a union of petals and the base of a vertex v, of eligible tenacity t, will be $\text{bud}^*(v)$ at the end of MAX in search level $(t-1)/2$.

EXAMPLE 8. In the graph of Figure 11, the two bridges (l_1, l_2) and (r_1, r_2) are of the same tenacity. The algorithm will break this tie arbitrarily and perform DDFS on these bridges in one of the two orders. In Figures 11 and 12, the order is (l_1, l_2) first and (r_1, r_2) second; these figures show the petals and buds found after DDFS is performed on the first and second bridge, respectively. Observe that b does not belong to the first petal but it does belong to the second petal. The blossom is the union of both petals and its base is f. The reader is encouraged to work out the petals if DDFS is performed on these bridges in the reverse order.

[13] To avoid cluttering up Figure 10, only two vertices are pointing to the petal-node.

5.3. The Mapping from Graph G to H We will give a succinct description of this mapping here; for a more in-depth treatment, see Sections 8.2.1 and 8.3.1. Each time DDFS is called, a new directed graph H is defined. It is a function of the bridge that triggers the current DDFS and the petals which have been found so far. A well-chosen subset of the vertices of G will form the vertices of H. For each vertex v of G that is chosen, its name in H will be v_H and we will define its level, $l(v_H) = \text{minlevel}(v)$.

Assume that DDFS is called with bridge (r, g). Then H must have the two vertices $\text{bud}^*(r)$ and $\text{bud}^*(g)$[14]. The rest of H is recursively defined as follows. If $\text{minlevel}(v) > 0$ then corresponding to each predecessor u of v in G, H has the vertex $\text{bud}^*(u)$ and the directed edge $(v, \text{bud}^*(u))$. If $\text{minlevel}(v) = 0$ then $l(v_H) = 0$ and v_H has no outgoing edges. It is easy to confirm that H satisfies the DDFS Requirement. For details see Sections 8.2.1 and 8.3.1.

Example 9. In the graph of Figure 13, DDFS called with bridge (u, v) ends in Case 2: the two centers of activity terminate at distinct unmatched vertices, f_1 and f_2. This indicates the presence of a minimum length augmenting path between f_1 and f_2. The next task is to find such a path.

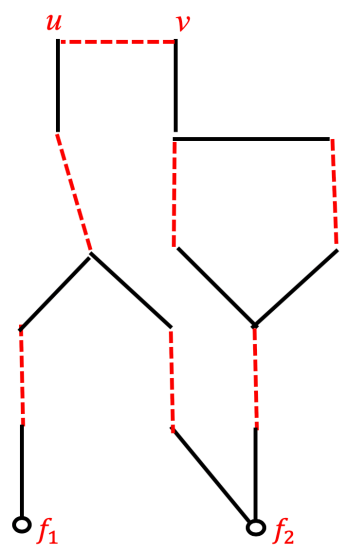

Figure 13. DDFS on the bridge (u, v) terminates with two unmatched vertices, f_1 and f_2, leading to an augmentation.

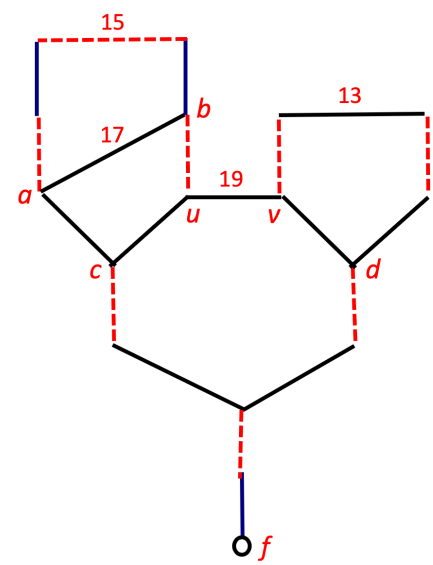

Figure 14. The tenacities of all four bridges is indicated. Bridge (a, b) has empty support.

Example 10. All bridges considered so far had non-empty support. However, this will not be the case in a typical graph, e.g., consider bridge (a, b) of tenacity 17, in Figure 14. Clearly the support of this bridge is $\emptyset$. DDFS will discover this right away since the two endpoints of this bridge have the same bud^*, i.e. $\text{bud}^*(a) = \text{bud}^*(b)$. Whether a bridge has empty support is not known a priori — it will become clear only after DDFS is performed on this bridge. Therefore, DDFS needs to be performed on every one of the bridges.

[14] It is possible that $\text{bud}^*(r) = \text{bud}^*(g)$. This happens if the bridge (r, g) has empty support and therefore a new petal is not formed, e.g., see Example 10.

5.4. Finding Augmenting Paths The MV algorithm will find augmenting paths during search level j_m, where $l_m = 2j_m + 1$ and l_m is the length of a minimum length augmenting path in the current phase. DDFS on some bridges of tenacity l_m will end in Case 2, i.e., no bottleneck is found. We note that not every bridge of tenacity l_m leads to an augmentation — DDFS on a bridge of tenacity l_m can end in Case 1 as well, i.e., a bottleneck is found.

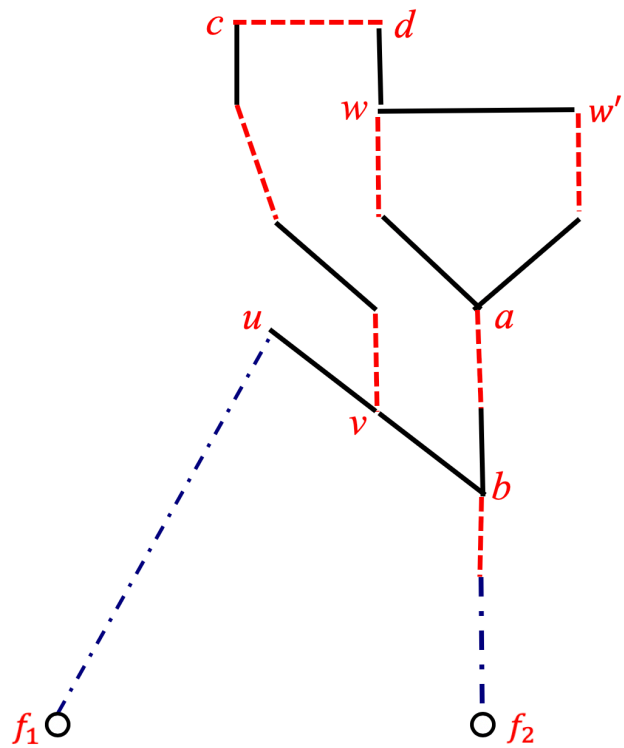

FIGURE 15. Constructing a minimum length augmenting path between unmatched vertices f_1 and f_2.

5.4.1. Finding one augmenting path In Figure 15, $\text{minlevel}(u) > \text{minlevel}(v)$ and therefore edge (u, v) is an anomalous bridge. When DDFS is performed on this bridge, assume that the red DFS trees has root u. Since v is already in a petal with $\text{bud}^*(v) = b$, the green DFS tree will have root b. The two trees will simply follow predecessors and will terminate at f_1 and f_2, respectively.

A DFS from u in the red tree will yield a path from u to f_1, say p_1. Since v is in a petal with $\text{bud}^*(v) = b$, the algorithm needs to find an alternating path from v to b, say p_2, starting with a matched edge, i.e., of even length. The construction of this path is described below[15]. Additionally, the algorithm needs to find a path, say p_3, from b to f_2 in the green tree of the DDFS performed on bridge (u, v). Then, the complete augmenting path from f_1 to f_2 will be $p_1^{-1} \circ (u, v) \circ p_2 \circ p_3$. Clearly, p_1 and p_3 are easy to find.

[15] In the language of Definition 23, p_2^{-1} is an $\text{evenlevel}(b; v)$ path.

We next describe how to find p_2. The algorithm observes that evenlevel$(v) =$ maxlevel(v) and therefore p_2 must use the bridge of the petal containing v. Using the petal node, the algorithm finds the endpoints of this bridge, namely c and d. It notices that c and v have the same color, say red. Therefore, it looks for a path from c to v in the red tree and a path from d to b in the green tree. For finding the latter path, it jumps from w to bud$^*(w)$ at the moment when DDFS was called with the bridge (c, d), since that will be the next node in the green tree. The latter node is a[16]. Next, the algorithm continues searching from a in the green tree and follows predecessors till it reaches b.

To find the complete path from d to b, the algorithm must find a path from a to w in the smaller petal[17]. This time, it observes that evenlevel$(w) =$ minlevel(w) and therefore the required path does not use the bridge of the smaller petal. Instead it is found by doing a DFS in the green tree, assuming that the color of w was green in the DDFS conducted on bridge (w, w'). Then p_2 is obtained by concatenating the path from v to c with (c, d) with the path from d to b. The latter consists of (d, w) concatenated with the path from w to a concatenated with the path from a to b.

Remark 1. Observe that the process of finding path p_2 from v to b, in Figure 15, critically used the markers left in the petal created by DDFS when called with bridge (c, d). In particular, the pointers from the petal-node to the endpoints of this bridge, i.e., c and d, are critical for ensuring linear running time: disjoint paths from c to v and d to b can be easily found by doing DFSs on the two trees rooted at c and d, see Definition 2. In the absence of these markers, the algorithm would need to find disjoint paths from v to c and b to d; it is unclear how this can be accomplished efficiently.

One last step in the process of finding an augmenting path, which could not be described since it requires Definition 27, is stated in Remark 6 in Section 8.3.1.

Example 11. In Figure 16, we have added one edge to the graph of Figure 5, namely (v, f') and the unmatched vertex f'. The evenlevel(v) path, which starts at f, followed by edge (v, f') yields the unique augmenting path in this figure. Therefore, finding evenlevel(v) in Figure 5 was not just an academic matter. However, the MV algorithm will not find this path by following the arrows in the figure; instead it will use the bridges and blossoms as detailed above. In particular, the bridge of tenacity $\gamma = 17$, as shown in Figure 7, triggers a DDFS which will find this path by skipping over the blossom of tenacity $\beta = 15$.

5.4.2. Finding a maximal set of disjoint paths After the first path, say p, is found, its vertices are removed. As a result, there may be other vertices that cannot be on minimum length augmenting paths that are disjoint from p. These vertices are recursively removed using the procedure RECURSIVE REMOVE which works as follows: First, all vertices of p and all edges incident at them are removed. If as a result there is a matched vertex v which has no more predecessors, it is removed. This process is continued until there are no more such vertices. Finally, all isolated unmatched vertices are removed.

At this point MAX will process the next bridge of tenacity l_m. When it encounters another bridge which makes DDFS terminate in Case 2, it finds another augmenting path. This continues until all bridges of tenacity l_m are processed. Lemma 24 shows that this will result in a maximal set of paths of length l_m.

6. Relationship between the Tenacity of an Edge and Tenacities of its Endpoints

The relationship depends on whether the edge is matched or unmatched, and in the latter case, whether it is a prop or a bridge. The answer in each case is significant and will influence the proof

[16] Observe that bud$^*(w)$ right after DDFS is performed on bridge (c, d) is b.

[17] Again, in the language of Definition 23, this is an evenlevel$(a; w)$ path.

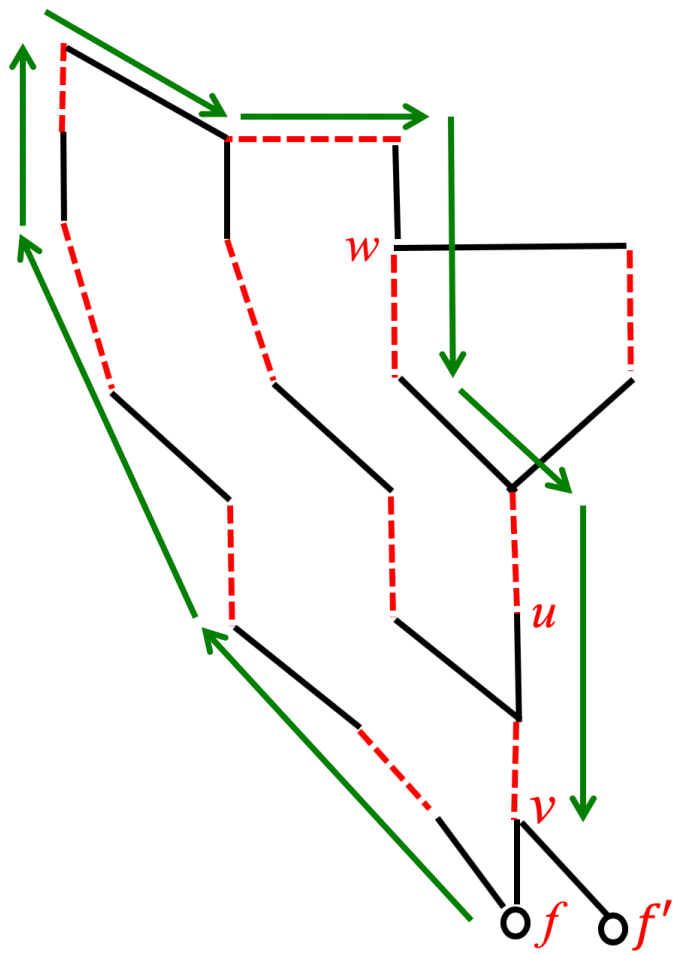

FIGURE 16. The evenlevel(v) path, which starts at f, followed by edge (v, f') is an augmenting path.

of the main theorem, Theorem 3, given in Sections 8.2 and 8.3. It will also help establish facts which give the correct way of synchronizing events in the algorithm, presented in Section 9.1.

LEMMA 1. *Let (u,v) be a matched edge of finite tenacity. Then* $\text{evenlevel}(v) = \text{oddlevel}(u) + 1$ *and* $\text{tenacity}(v) = \text{tenacity}(u,v)$. *Furthermore, (u,v) is a bridge if and only if* $\text{oddlevel}(u) = \text{oddlevel}(v) = i$, *where* $\text{tenacity}(u,v) = 2i+1$.

Proof : Let p be an $\text{oddlevel}(u)$ path. If v lies on p then (u,v) must also lie on p, since p is an alternating path. If so, p is not a simple odd length alternating path to u, giving a contradiction. Therefore, v does not lie on p and hence $p \circ (u,v)$ is a minimum even length alternating path to v. Therefore, $\text{evenlevel}(v) = \text{oddlevel}(u) + 1$ and similarly $\text{evenlevel}(u) = \text{oddlevel}(v) + 1$. Therefore $\text{tenacity}(v) = \text{evenlevel}(v) + \text{oddlevel}(v) = \text{oddlevel}(u) + 1 + \text{oddlevel}(v) = \text{tenacity}(u,v) = \text{evenlevel}(u) + \text{oddlevel}(u) = \text{tenacity}(u)$.

Now, there are two cases, either one of u or v has an oddlevel of $< i$ or $\text{oddlevel}(u) = \text{oddlevel}(v) = i$. In the first case, assume $\text{oddlevel}(u) < i$. If so, u is a predecessor of v and (u,v) is a prop. In the second case, neither endpoint is a predecessor of the other and (u,v) is a bridge. Finally, if (u,v) is a bridge, the first case cannot apply. Therefore $\text{oddlevel}(u) = \text{oddlevel}(v) = i$. □

REMARK 2. As a consequence of Lemma 1, in several proofs, restricting attention to only one of the end points of a matched edge (u,v) will suffice, since $\text{evenlevel}(v)$ and $\text{oddlevel}(v)$ fully determine the two levels of u.

LEMMA 2. *Let (u,v) be an unmatched edge of finite tenacity and assume that u is a predecessor of v. Then* $\text{tenacity}(v) = \text{tenacity}(u,v)$.

Proof : We will show that $\text{oddlevel}(v) = \text{evenlevel}(u) + 1$. If so, $\text{tenacity}(v) = \text{evenlevel}(v) + \text{oddlevel}(v) = \text{evenlevel}(v) + \text{evenlevel}(u) + 1 = \text{tenacity}(u,v)$, thereby proving the lemma. Since v is getting its minlevel from the unmatched edge (u,v), $\text{minlevel}(v) = \text{oddlevel}(v)$ and therefore v is an inner vertex.

Since u is a predecessor of v, there is an $\text{oddlevel}(v)$ path, say q, that ends with the edge (u,v). Now, $\text{oddlevel}(v) = |q|$ and $\text{evenlevel}(u) \le |q| - 1$. Therefore, $\text{oddlevel}(v) > \text{evenlevel}(u)$.

Let p be an $\text{evenlevel}(u)$ path. First assume that v lies on p. If v is odd w.r.t. p[18], then $\text{oddlevel}(v) < \text{evenlevel}(u)$, a contradiction. Therefore, v is even w.r.t. p. But then $\text{evenlevel}(v) < \text{evenlevel}(u)$ implying that $\text{evenlevel}(v) < \text{oddlevel}(v)$ and that v is an outer vertex, another contradiction. Therefore v does not lie on p. Therefore $p \circ (u,v)$ is an $\text{oddlevel}(v)$ path and therefore $\text{oddlevel}(v) = \text{evenlevel}(u) + 1$ holds, giving the lemma. □

LEMMA 3. *Let (u,v) be an unmatched bridge. Then:*
1. $\text{tenacity}(v) \le \text{tenacity}(u,v)$.
2. *If* $\text{tenacity}(v) = \text{tenacity}(u,v)$ *then v is an outer vertex.*

Proof : 1). We will show that $\text{oddlevel}(v) \le \text{evenlevel}(u) + 1$. If so, adding $\text{evenlevel}(v)$ to both sides of this inequality we will get $\text{tenacity}(v) \le \text{tenacity}(u,v)$.

Let p be an $\text{evenlevel}(u)$ path starting at unmatched vertex f, say. If v lies on p, there are two cases. If v is odd w.r.t. p, then $p[f \text{ to } v]$ is an odd length alternating path from f to v and therefore $\text{oddlevel}(v) \le |p[f \text{ to } v]| < |p| = \text{evenlevel}(u)$.

Next assume that v is even w.r.t. p. Then $p[f \text{ to } v] \circ (v,u)$ is an odd length alternating path from f to u. Clearly the length of this path is less than that of p, implying that u is an inner vertex. If this path were an $\text{oddlevel}(u)$ path, then v would be a predecessor of u, implying that (u,v) is a prop and contradicting the fact that (u,v) is a bridge. Therefore $\text{oddlevel}(u) < |p[f \text{ to } v]| + 1$. Let q be an $\text{oddlevel}(u)$ path which starts at unmatched vertex f', say, where $f = f'$ is allowed. Let w be the first vertex of q that lies on $p[v \text{ to } u]$. If w is odd w.r.t. p then $q[f' \text{ to } w] \circ p[w \text{ to } u]$ is an even length alternating path from f' to u of length less than $|p| = \text{evenlevel}(u)$, giving a contradiction. Therefore w is even w.r.t. p. If so, $q[f' \text{ to } w] \circ p[w \text{ to } v]$ is an odd length alternating path from f' to v of length less than $|p|$. Therefore $\text{oddlevel}(v) < \text{evenlevel}(u)$.

Next, assume that v does not lie on p. Then, $p \circ (u,v)$ is an odd length alternating path from f to v. Therefore $\text{oddlevel}(v) \le \text{evenlevel}(u) + 1$.

2). Assume that $\text{tenacity}(v) = \text{tenacity}(u,v)$. Then $\text{oddlevel}(v) = \text{evenlevel}(u) + 1$. If v were an inner vertex, then $\text{oddlevel}(v) = \text{minlevel}(v)$ and the previous equality implies that u is a predecessor of v and (u,v) is a prop, leading to a contradiction. Therefore v is an outer vertex. □

REMARK 3. Let (u,v) be an unmatched edge and let u be a predecessor of v. Then $\text{tenacity}(u)$ can be smaller than, equal to, or bigger than $\text{tenacity}(u,v)$. The first case is illustrated by prop (r_1, l_2) and the third case by the props out of f in Figure 10. Let (u,v) be an unmatched bridge and assume that $\text{tenacity}(v) < \text{tenacity}(u,v)$. Then v can be an inner or an outer vertex. The bridge (a,b) in Figure 14 illustrates both possibilities; although (a,b) has empty support, it is easy to show this for bridges with non-empty support as well.

[18] For this notion, see Definition 7.

7. Limited BFS-Honesty Consider a minimum length alternating path, p, from unmatched vertex f to a vertex v; p is allowed to be of either parity. The notion of tenacity enables us to characterize a subset of vertices of p that are BFS honest on p, namely all vertices on p whose tenacity is at least as large as that of v. This BFS-honesty will be critically exploited later.

Definition 15. (even/odd w.r.t. p) Let p be an evenlevel(v) or oddlevel(v) path starting at unmatched vertex f and let u lie on p. Then $|p[f \text{ to } u]|$ will denote the length of this path from f to u, and if it is even (odd) we will say that u is *even (odd) w.r.t. p*.

Definition 16. (BFS honesty on p) Let p be an evenlevel(v) or oddlevel(v) path starting at unmatched vertex f and let u lie on p. We will say that u is *BFS honest on p* if $|p[f \text{ to } u]| =$ evenlevel(u) (oddlevel(u)) if u is even (odd) w.r.t. p.

Example 12. Observe that the graphs of Figures 5 and 6 are identical, with vertex names given in the former and vertex tenacities in the latter. The vertices u and v are BFS honest on all evenlevel and oddlevel paths to the vertices of tenacity α. However, the two vertices of tenacity α that lie on the evenlevel(u) and evenlevel(v) paths are not BFS honest on these paths.

Theorem 2. *Let p be an* evenlevel(v) *or* oddlevel(v) *path starting at unmatched vertex f and let vertex $u \in p$ with* tenacity(u) $\geq$ tenacity(v). *Then u is BFS honest on p. Furthermore, if* tenacity(u) $>$ tenacity(v) *then* $|p[f \text{ to } u]| =$ minlevel(u).

Proof : Assume w.l.o.g. that p is an evenlevel(v) path and that u is even w.r.t. p (by Lemma 1). Suppose u is not BFS honest on p, and let q be an evenlevel(u) path, i.e., $|q| < |p[f \text{ to } u]|$. First consider the case that evenlevel(v) = maxlevel(v), and let r be a minlevel(v) path. Let u' be the matched neighbor of u. Consider the first vertex of r that lies on $p[u' \text{ to } v]$. If this vertex is even w.r.t. p then oddlevel(u) $\leq |r| + |p[u \text{ to } v]|$. Additionally, evenlevel(u) $< |p[f \text{ to } u]|$, hence tenacity(u) $<$ tenacity(v), leading to a contradiction. On the other hand, if this vertex is odd w.r.t. p then minlevel(v) $= |r| >$ evenlevel(u), because otherwise there is a shorter even path from f to v than evenlevel(v). We combine the remaining argument along with the case that evenlevel(v) = minlevel(v) below.

Consider the first vertex, say w, of q that lies on $p(u \text{ to } v]$ – there must be such a vertex because otherwise there is a shorter even path from f to v than evenlevel(v). If w is odd w.r.t. p then we get an even path to v that is shorter than evenlevel(v). Hence w must be even w.r.t. p. Then, $q[f \text{ to } w] \circ p[w \text{ to } u]$ is an odd path to u with length less than evenlevel(v), where $\circ$ denotes the concatenation operator. Again we get tenacity(u) $<$ tenacity(v), leading to a contradiction.

We next prove the second claim of the theorem. First consider the case that evenlevel(v) = minlevel(v), and assume for contradiction that $|p[f \text{ to } u]| =$ maxlevel(u). Then tenacity(u) $< 2 \cdot$ maxlevel(u) $< 2 \cdot |p[f \text{ to } v]| = 2 \cdot$ minlevel(v) $<$ tenacity(v), a contradiction.

Therefore, evenlevel(v) = maxlevel(v). As before, let r be a minlevel(v) path, and consider the first vertex of r that lies on $p[u \text{ to } v]$. If this vertex is even w.r.t. p then oddlevel(u) $\leq |r| + |p[u \text{ to } v]|$. Hence tenacity($u$) $\leq$ tenacity(v), which leads to a contradiction. On the other hand, if this vertex is odd w.r.t. p then minlevel(v) $= |r| >$ evenlevel(u), because otherwise there is a shorter even path from f to v than evenlevel(v). Now the claim follows because otherwise tenacity(u) $<$ tenacity(v). □

Corollary 1. *Let p be an* evenlevel(v) *or* oddlevel(v) *path and let u lie on p. If u is not BFS honest on p then* tenacity(u) $<$ tenacity(v).

8. Base, Blossom and Bridge Recall Definitions 3 and 9 which defined l_m and t_m as the length of a minimum length augmenting path and the tenacity of a minimum tenacity vertex, respectively. As noted earlier, $t_m \leq l_m$ and the case $t_m = l_m$ is trivial. For the rest of the paper, we will deal with the main case, namely $t_m < l_m$; Example 14 explains the importance of this

assumption. In Definition 17, we introduce the notion of eligible tenacity. Theorem 3 helps establish the central notions of base and blossom for vertices of eligible tenacity.

Definition 17. (Eligible tenacity) An odd number t, with $t_m \le t < l_m$, will be said to be an *eligible tenacity.*

Definition 18. (Higher and lower on a path) Let v be a vertex and p be an evenlevel(v) or oddlevel(v) path; assume it starts at unmatched vertex f. If u and w are two vertices on p and if u is further away from f on p than w, then we will say that u is *higher* than w and w is *lower* than u on p.

Definition 19. (The set $B(v)$) Let v be a vertex of eligible tenacity and p be an evenlevel(v) or oddlevel(v) path starting at unmatched vertex f. Let $t = \text{tenacity}(v)$ and consider all vertices of tenacity $> t$ on p; clearly, this set is nonempty since contains f. Among these vertices, define the highest one to be *the base of v w.r.t. p*, denoted by $F(p,v)$. Clearly, $F(p,v)$ is even w.r.t. p, and by Theorem 2, it is an outer vertex. Finally define:

$$B(v) = \{F(p,v) \mid p \text{ is an evenlevel}(v) \text{ or oddlevel}(v) \text{ path}\}.$$

8.1. Central Structural Facts As mentioned in Section 1.1, the central structural fact needed to prove correctness of the algorithm is the following:

∗ *For every vertex v such that* $\text{tenacity}(v) = t$ *and* $t_m \le t \le l_m$*, every* $\text{maxlevel}(v)$ *path contains a bridge of tenacity t.*

The proof of this fact requires several other structural facts. Among them, the most important one is the following:

∗ *For a vertex v of eligible tenacity, the set $B(v)$ is a singleton.*

Once this fact is proven, we can define the *base of v* to be the vertex in $B(v)$ and move on to defining the notion of a blossom and proving its properties. However, the proof of this fact is not straightforward because of the following "chicken and egg problem", which was also mentioned in Section 1.1: On the one hand, the proof of this fact requires the notion of blossom and its associated properties, and on the other, blossoms can be defined only after defining the base of a vertex.

We will break this deadlock by proving this fact via an induction on eligible tenacities. Once this fact is proven for vertices of tenacity $\le t$, the base of vertices of tenacity $\le t$ can be defined. Following this, blossoms of tenacity t can be defined and properties of these blossoms, and properties of paths traversing through these blossoms, can be established. These properties are then used to prove this fact for tenacity $t+2$.

8.2. The Induction Basis This section is devoted to proving the induction basis for Theorem 3. This involves proving all statements mentioned in Theorem 3 for the case $t = t_m$; each statement is proven in a separate lemma. For this purpose we will define a subgraph of G, namely H'_m. Its structure is fairly simple, thereby making the proofs easy; in contrast, the analogous graph for the induction step is considerably more complex. The saving grace is that the proofs for the base case provide much insight on how to proceed with the induction step.

Lemma 4. *Let (u,v) be an edge of tenacity t_m. Then (u,v) is a bridge if and only if* $\text{minlevel}(u) = \text{minlevel}(v) = i$*, where* $t_m = 2i+1$.

Proof : If (u,v) is matched, the claim follows by Lemma 1. Next assume that (u,v) is unmatched. If so, $\text{tenacity}(u,v) = \text{evenlevel}(u) + \text{evenlevel}(v) + 1$. First assume (u,v) is a bridge. Since the tenacity of a vertex cannot be less than t_m, by Lemma 3, $\text{tenacity}(u) = \text{tenacity}(v) = \text{tenacity}(u,v)$, and u and v are both outer vertices. Therefore, $\text{minlevel}(u) = \text{evenlevel}(u)$ and $\text{minlevel}(v) = \text{evenlevel}(v)$. Therefore, $\text{oddlevel}(v) = \text{evenlevel}(u) + 1$ and $\text{oddlevel}(u) = \text{evenlevel}(v) + 1$.

Since $t_m = 2i+1$, $\text{minlevel}(u) = \text{evenlevel}(u) \le i$ and $\text{minlevel}(v) = \text{evenlevel}(v) \le i$. Assume one of them has minlevel $< i$, say u. Then, $\text{oddlevel}(v) = \text{evenlevel}(u) + 1 \le i$. This implies that

evenlevel$(v) < i$, since v is outer, giving tenacity$(v) < 2i$, a contradiction. Therefore, minlevel$(u) =$ minlevel$(v) = i$.

Next, assume minlevel$(u) =$ minlevel$(v) = i$. Since (u, v) is an unmatched edge of tenacity t_m, evenlevel$(u) +$ evenlevel$(v) + 1 = t_m$. Therefore evenlevel$(u) =$ evenlevel$(v) = i$. Now, since (u, v) unmatched, u is not a predecessor of v and v is not a predecessor of u. Therefore (u, v) is a bridge. □

LEMMA 5. *Let v be a vertex of tenacity t_m and p be a* maxlevel(v) *path. Then p contains a unique bridge of tenacity t_m.*

Proof : Assume that p starts at unmatched vertex f. Let q be a minlevel(v) path and assume it starts at unmatched vertex f', which may or may not be the same as f. If p and q meet only at v then $f \neq f'$ and $p \circ q$ is an augmenting path of length t_m, leading to a contradiction. Therefore the intersection of p and q contains vertices in addition to v. Let u be the highest vertex of $p[f$ to $v)$ that is also on q.

Now there are two cases. If $u = f$, then $f = f'$ and p and q meet at two vertices, namely $u = f$ and v. Next assume that $u \neq f$. If so, u is matched and p must contain the matched edge incident at u. Therefore the vertex u must be even w.r.t. p. By definition, tenacity$(u) \geq t_m$, and therefore by Theorem 2, u is BFS honest on p as well as q. Therefore, evenlevel$(u) = |p[f$ to $u]|$. If u is odd w.r.t. q, then oddlevel$(u) = |q[f'$ to $u]|$ thereby implying that tenacity$(u) < |p \circ q| = t_m$, a contradiction. Therefore u is even w.r.t. q as well and $|q[f'$ to $u]| =$ evenlevel(u).

Therefore in both cases, $p[u$ to $v] \circ q[v$ to $u]$ is an odd length cycle having two unmatched edges incident at u and is fully matched otherwise. By concatenating $p[f$ to $u]$ to an appropriate subpath of this cycle, we can obtain even and odd alternating paths to all vertices of this cycle, other than u. For any vertex $w \neq u$ on this cycle, the sum of the lengths of the even and odd paths is t_m. Therefore, these must be minimum length alternating paths and tenacity$(w) = t_m$. Furthermore, since this cycle has length at least 3, evenlevel$(u) = |p[f$ to $u]| < i$.

Assume that this odd cycle has length $2k+1$ and number its edges consecutively, starting at u. Let (w, w') be the $k+1^{st}$ edge, i.e., the middle edge. Then clearly, minlevel$(w) =$ minlevel$(w') = i$. Therefore, by Lemma 4, (w, w') is a bridge of tenacity t_m. Clearly, besides w and w', no other vertices on p can have minlevel of i. Therefore there are no other bridges of tenacity t_m on it. □

Note that in the proof given above, $v \in$ support(w, w'). In general, v may lie in the support of several bridges of tenacity t_m.

8.2.1. The Graphs H_m and H'_m Proving that the set $B(v)$ is a singleton is not straightforward even for vertices of tenacity t_m. A major simplification is achieved by using the power of DDFS — in particular, the DDFS Certificate provided on its termination. DDFS is carried out on a special directed, layered graph H_m, which satisfies the DDFS Requirement. We start by defining H_m and a closely related graph, H'_m; the latter is a subgraph of G.

Let $U_m = \{v \in V \mid \text{tenacity}(v) = t_m\}$. Let B_m denote the set of all bridges of tenacity t_m and let T_m denote the set of end points of bridges in B_m.

LEMMA 6. *Let $v \in T_m$ and let p be a* minlevel(v) *path. Then all edges on p are props.*

Proof : Assume that p starts at unmatched vertex f and let u be any vertex on p other than v. Clearly tenacity$(u) \geq$ tenacity(v) and therefore by Theorem 2, u is BFS honest on p. Furthermore $|p[f$ to $u]| =$ minlevel(u), since otherwise tenacity$(u) <$ tenacity(v). Therefore p gives minlevels to all vertices on it and hence all its edges are props. □

Consider each vertex $v \in T_m$ and let p denote an arbitrary minlevel(v) path. Denote by V_m the union of all vertices on all such paths p for all vertices $v \in T_m$. Furthermore, denote by P_m the

union of all edges on all such paths p. By Lemma 6, all edges in P_m are props in the original graph G. Define $H'_m = (V_m, (P_m \cup B_m))$ and define a matching in H'_m as follows: edge $e \in (P_m \cup B_m)$ is matched if and only if e is matched in G. Clearly, H'_m is a subgraph of G.

The next definition is related to Definition 11.

Definition 20. (The relations pred_m and pred^*_m) If $u \in U_m$, $v \in V_m$ and v is a predecessor of u, then we will say that v *is* pred_m *of* u, denoted by $v = \text{pred}_m(u)$. Observe that under this definition, we do not allow $\text{tenacity}(u) > t_m$; $\text{tenacity}(u)$ *must be* t_m. On the other hand, $\text{tenacity}(v) > t_m$ is allowed and v may even be an unmatched vertex. Next let $u \in U_m$ and $v \in V_m$, with $u \neq v$, such that there is a path from u to v in H'_m and all vertices on this path are in U_m, except possibly v. Then we will say that v *is* pred^*_m *of* u, denoted by $v = \text{pred}^*_m(u)$.

Lemma 7. *For every $v \in U_m$, H'_m contains all possible* $\text{evenlevel}(v)$ *and* $\text{oddlevel}(v)$ *paths that are present in G.*

Proof : We will use the fact that, by construction, H'_m contains all possible $\text{minlevel}(u)$ and $\text{maxlevel}(u)$ paths for all vertices $u \in T_m$.

For $v \in U_m$, let p and q be $\text{minlevel}(v)$ and $\text{maxlevel}(v)$ paths, respectively, in G, which start at unmatched vertex f, say. Let (u, u') be the unique bridge of tenacity t_m on q, with u' being higher than u on q; for definition of "higher on a path" see Definition 18. Then, $q[f \text{ to } u]$ is a $\text{minlevel}(u)$ path and therefore it is present in H'_m. Furthermore, $p \circ q[v \text{ to } u']$ is a $\text{minlevel}(u')$ path and is also present in H'_m. Therefore H'_m contains p and q, hence proving the lemma. □

Next we will define the directed, layered graph H_m. For each edge $(u, v) \in P_m$, direct it from v to u if u is a predecessor of v. To keep notation simple, we will denote this set of directed edges by P_m as well; the context will easily clarify which graph is being referred to.

We will partition V_m into $i+1$ layers numbered $0, \ldots, i$, where $t_m = 2i+1$ and the layer number of $v \in V_m$ is $l(v) = \text{minlevel}(v)$. Let $H_m = (V_m, P_m)$. Observe that each edge of H_m runs from a layer, say l, to $l-1$, i.e., it is of unit length. In contrast, the graph H_t, defined in Section 8.3.1, has long edges. It is easy to check that H_m satisfies the DDFS Requirement. We will use the information obtained from DDFS on H_m to find minlevel and maxlevel paths in H'_m. The mapping between the vertices of H_m and H'_m is the obvious one.

Corresponding to each bridge $(u, v) \in B_m$, conduct DDFS in H_m starting at u and v, and denote by $k_{(u,v)}$ the bottleneck found. Clearly, each bottleneck is an outer vertex. Note that two different bridges may have the same bottleneck. Using the DDFS Certificate and the mapping between H_m and H'_m we get:

Lemma 8. *Let bridge $(u, v) \in B_m$ and let $w \in \text{support}(u, v)$. Then:*
1. *$w = \text{pred}^*_m(u)$ or $w = \text{pred}^*_m(v)$ or both.*
2. *$k_{(u,v)} = \text{pred}^*_m(w)$.*

Procedure Bottleneck: The input to this procedure is a bridge $(u, v) \in B_m$.
Conduct DDFS in H_m, starting at u and v, to find the bottleneck $k_{(u,v)}$.
If $\text{tenacity}(k_{(u,v)}) > t_m$, HALT.
Otherwise, there is a bridge of tenacity t_m, say (u', v'), such that $k_{(u,v)} \in \text{support}(u', v')$. Conduct DDFS on bridge (u', v') to find its bottleneck, $k_{(u',v')}$. Clearly, $\text{minlevel}(k_{(u',v')}) < \text{minlevel}(k_{(u,v)})$, and $k_{(u',v')} = \text{pred}^*_m(k_{(u,v)})$. If $\text{tenacity}(k_{(u',v')}) = t_m$ repeat this process until a bottleneck, say b, is encountered which has tenacity $> t_m$. This is bound to happen because the minlevels of the bottlenecks are decreasing and eventually the bottleneck will turn out to be an unmatched vertex, say f, and f satisfies $\text{tenacity}(f) \geq l_m > t_m$.

Example 13. Let us illustrate Procedure Bottleneck on the graph of Figure 18. When called with bridge (u, u'), the bottleneck found is b'. However, $\text{tenacity}(b') = t_m = 15$ and b' is in the

support of (v, v'). DDFS on this bridge will result in the bottleneck of b. Since tenacity$(b) > 15$, the procedure halts and b is the base of the endpoints of both bridges.

The vertex b identified by Procedure Bottleneck is very special, as will be established next. It is called a *base*; a formal definition is given below. Clearly, b is an outer vertex. In general, H'_m will have a number of bases.

For each base b in H'_m, define the set

$$S_{b,t_m} = \{v \in U_m \mid b \text{ is pred}^*_m \text{ of } v\}.$$

Observe that $b \notin S_{b,t_m}$ and if b, b' are two bases in H'_m then $S_{b,t_m} \cap S_{b',t_m} = \emptyset$. The next lemma is implied by Lemma 1 and the fact that set S_{b,t_m} is defined via Procedure Bottleneck.

LEMMA 9. *Let $v \in S_{b,t_m}$ and let v' be the matched neighbor of v. Then $v' \in S_{b,t_m}$.*

LEMMA 10. *Let $v \in S_{b,t_m}$. Then every* evenlevel(v) *and* oddlevel(v) *path in the graphs H'_m and G consists of an* evenlevel(b) *path followed by an alternating path using vertices of S_{b,t_m}.*

Proof : Let $(u, u') \in B_m$ be a bridge of tenacity t_m whose endpoints are in S_{b,t_m}, and let p be any minlevel(u) path. By Lemma 8 and the procedure given above, p^{-1} starts at u and follows down predecessors until it arrives at $k_{(u,u')}$. If $k_{(u,u')} \neq b$, p^{-1} follows down predecessors until it arrives at b. In either case, the rest of p^{-1} is an evenlevel(b) path followed in reverse. Therefore, p has the structure described in the statement of the lemma.

Next, assume that $v \in S_{b,t_m}$ and let p and q be minlevel(v) and maxlevel(v) paths in G. Using the arguments given in Lemma 7, $q \circ p^{-1}$ can be decomposed into minlevel(u) and minlevel(v) paths and the unique bridge on q. Now, by the assertion made above, p and q also have the structure described in the statement of the lemma. □

COROLLARY 2. *For every $v \in U_m$, the set $B(v)$ is a singleton.*

DEFINITION 21. (The base of a vertex of tenacity t_m and basal vertices) For each $v \in U_m$, define base(v) to be the unique vertex, say b, in the set $B(v)$. We will say that the *base* of v is b. Each such vertex b will be called a *basal vertex*. Clearly, b is an outer vertex and tenacity$(b) > t_m$.

EXAMPLE 14. In the graph of Figure 17, vertices u and v do not have a base, even though they are of finite tenacity. Clearly tenacity(u) = tenacity(v) = tenacity(f_1) = tenacity$(f_2) = 3$ and the graph has no vertex of tenacity greater than 3. Since $l_m = 3$, u and v are not of eligible tenacity; this explains why they do not have a base. The following question arises: can f_1 be viewed as the base of u and v, even though its tenacity is the same as that of u and v? A negative answer is easy to see after adding unmatched edge (f_2, v), since then the graph has minlevel(v) and maxlevel(u) paths which don't contain f_1.

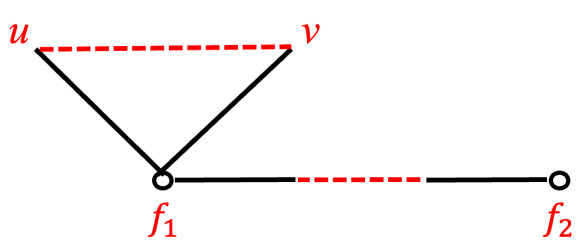

FIGURE 17. Vertices u and v have no base. The tenacity of bridge (u, v) is 3 and so is the tenacity of u, v, f_1 and f_2.

DEFINITION 22. (Blossom of tenacity t_m and base b) Let b be a basal vertex as defined in Definition 21. Then the *blossom of tenacity t_m and base b* is the set $\mathcal{B}_{b,t_m} = \{v \in U_m \mid \text{base}(v) = b\}$.

Clearly, $\mathcal{B}_{b,t_m} \neq \emptyset$, $\mathcal{B}_{b,t_m} = S_{b,t_m}$, $b \notin \mathcal{B}_{b,t_m}$ and for each vertex $v \in \mathcal{B}_{b,t_m}$, $b = \text{pred}^*_m(v)$. The next fact follows from Lemmas 1 and 9.

Corollary 3. *Let (u,v) be a matched edge of tenacity t_m. Then u and v have the same base, say b, and both belong to the same blossom of tenacity t_m, namely $\mathcal{B}_{b,t_m}$.*

Example 15. Blossoms of tenacity t_m can be quite complex, as illustrated in Figure 18, even though they do not contain nested blossoms. This blossom has two bridges of tenacity t_m, both having non-empty support, and therefore it will have two petals, see Section 5.2 for this notion. The exact petals will depend on the order in which DDFS is conducted on bridges (u,u') and (v,v').

Definition 23. (Shortest path from a base to a vertex of tenacity t_m) Let $v \in U_m$ and $b = \text{base}(v)$. Then by an evenlevel$(b;v)$ (oddlevel$(b;v)$) path we mean a minimum even (odd) length alternating path in G from b to v that starts with an unmatched edge.

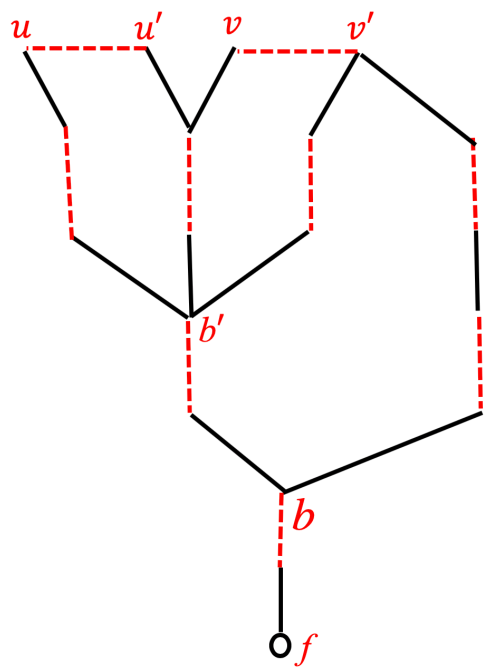

Figure 18. The set of vertices having base b form a blossom of tenacity 15 and $t_m = 15$. Observe that b' is not basal. This blossom contains the endpoints of bridges (u,u') and (v,v').

Lemma 11. *Let $v \in U_m$ and $b = \text{base}(v)$. Let p be an* evenlevel(b) *path and q be an* evenlevel$(b;v)$ *or* oddlevel$(b;v)$ *path. Then q meets p at b only.*

Proof : By Lemma 1, it suffices to prove this lemma for an oddlevel$(b;v)$ path q. For contradiction, assume that q meets p at vertices besides b. By Lemma 10, oddlevel$(v) \geq |p| + |q|$. We will define certain subpaths of q as *segments* as follows. Follow along q from b until it meets p, at w say. Then $q[b \text{ to } w]$ will be called a segment. Subsequent to this, each time q leaves p, at vertex r, say, and meets up p again, at s, say, then $q[r \text{ to } s]$ is called a segment. Eventually, q leaves p at y, say, and ends up at v. Then, $q[y \text{ to } v]$ is the last segment.

Now, there are two cases: either there is a segment, say $q[u \text{ to } w]$ such that u and w are both outer vertices or there is no such segment. In the first case, consider the odd length cycle $q[u \text{ to } w] \circ$

$p[w \text{ to } u]$ and assume that u is higher than w on p. Then, all vertices on this cycle, other than w have tenacity $< t_m$, a contradiction.

In the second case, since the first segment starts at an outer vertex, namely b, it ends at an inner vertex. Therefore, the next segment again starts at an outer vertex and so on. Finally, for the last segment, $q[y \text{ to } v]$, y must be an outer vertex. Assume that p starts at unmatched vertex f. If so, $p[f \text{ to } y] \circ q[y \text{ to } v]$ is a shorter odd alternating path from f to v than $|p| + |q|$, leading to another contradiction. The lemma follows. □

Lemmas 10 and 11 give:

Corollary 4. *Let $v \in U_m$ and $b = \text{base}(v)$. Then the matched neighbor of v also lies in the blossom and every* evenlevel(v) (oddlevel(v)) *path consists of an* evenlevel(b) *path followed by an* evenlevel$(b; v)$ (oddlevel$(b; v)$) *path, where the latter lies in $\mathcal{B}_{b,t_m} \cup \{b\}$.*

Lemma 10 and Corollary 4 give:

Lemma 12. *Let $\mathcal{B}_{b,t_m}$ and $\mathcal{B}_{b',t_m}$ be two blossoms with bases $b \neq b'$. Then $\mathcal{B}_{b,t_m} \cap \mathcal{B}_{b',t_m} = \emptyset$.*

Finally we need to prove one more fact, which captures the disciplined manner in which a minimum length alternating path, which intersects a blossom of tenacity t_m, $\mathcal{B}_{b,t_m}$, uses vertices of $(\mathcal{B}_{b,t_m} \cup \{b\})$. This fact will be used in the induction step.

Lemma 13. *Let p be an* evenlevel(v) *path from unmatched vertex f to an arbitrary vertex v such that there is a blossom $\mathcal{B}_{b,t_m}$ with $p \cap \mathcal{B}_{b,t_m} \neq \emptyset$. Then the base of this blossom, b, also lies on p and there is a vertex $u \in (p \cap \mathcal{B}_{b,t_m})$ such that $p[b \text{ to } u]$ contains all the vertices of $p \cap (\mathcal{B}_{b,t_m} \cup \{b\})$ and $p[b \text{ to } u]$ is an* evenlevel$(b; u)$ *path.*

Proof : For the sake of contradiction, assume that p does not intersect $\mathcal{B}_{b,t_m} \cup \{b\}$ in the manner described above. If so, two cases arise.

Case 1: b lies on p and there is a vertex $u \in (p \cap \mathcal{B}_{b,t_m})$ such that $p[b \text{ to } u]$ contains all the vertices of $p \cap (\mathcal{B}_{b,t_m} \cup \{b\})$ as well as some vertices not in $\mathcal{B}_{b,t_m}$, i.e., $p[b \text{ to } u]$ enters and exits $(\mathcal{B}_{b,t_m} \cup \{b\})$ more than once.

If so, by Corollary 4, $p[b \text{ to } u]$ is longer than an evenlevel$(b; u)$ path. Therefore, replacing the former with the latter will result in a shorter even path to v, giving a contradiction.

Case 2: The first vertex of p in $(\mathcal{B}_{b,t_m} \cup \{b\})$ is $x \in \mathcal{B}_{b,t_m}$, the last vertex is $y \in \mathcal{B}_{b,t_m}$ and $p(x \text{ to } y)$ is arbitrary, i.e., it may visit b as well as other vertices not in $\mathcal{B}_{b,t_m}$. For ease, we will call $p[f \text{ to } x]$ the *blue path* and $p[y \text{ to } v]$ the *red path*.

Let q be an evenlevel(b) path starting at unmatched vertex f', say, where $f' = f$ is allowed. Now there are several cases. If the red path does not intersect q then by Corollary 4, $q \circ s \circ p[y \text{ to } v]$ is a shorter even path[19] to v than p, where s is an evenlevel$(b; y)$ path.

If the red path does intersect q and the last vertex of p on q is even w.r.t. q, say w, then follow q to w and then follow by p to v; this is shorter than p. Next assume that the last vertex of p on q is odd w.r.t. q, say w. Then we ask whether the blue path intersects $q(w \text{ to } b]$. If the answer is no, then via Corollary 4 we get that the blue path, followed by r^{-1} to b, followed by q to w, followed by p to v is a shorter even path to v, where r is an evenlevel$(b; x)$ path.

The main case: Finally, assume that the blue and red paths both intersect q in arbitrary ways and the last vertex of the red path on q is odd w.r.t. q, say w. Both paths will traverse some matched edges of the path q. We will say that such an edge is blue (red) if the blue (red) path

[19] For clarity, in subsequent cases we will describe the path constructed in plain English as follows: Follow q to b, then s to y, then the red path to v, where s is an evenlevel$(b; y)$ path.

traverses it; furthermore, we will direct the edge in the direction in which the path traverses it. We will say that such an edge is *directed up* if it points towards b and is *directed down* if it points towards f'.

A subpath of $q[w \text{ to } b)$ is said to be *unicolored* if it does not contain matched edges of both colors; it is allowed to contain uncolored matched edges. A maximal unicolored subpath which starts and ends with colored matched edges is called an *interval*. In each interval, we will *mark* one edge as follows: in a blue interval, mark the lowest[20] edge on the blue path and in a red interval, mark the highest edge on the red path. The rest of the arguments will only be based on the marked edges and not the rest of the intervals. Clearly on $q[w \text{ to } b)$, the color of the marked edges must be alternating, with the lowest marked edge being red and directed down.

There are two cases: either all marked red edges are directed down or not. In the latter case, let e_1 be the lowest marked red edge which is directed up and let e_2 be the marked red edge just below e_1; e_2 is directed down. Let e_3 be the marked blue edge in between e_1 and e_2. Now if e_3 is directed down then follow the blue path from f to e_3, then follow q to e_2 and finally follow the red path all the way to v. However, if e_3 is directed up then follow the blue path from f to e_3, then follow q to e_1 and finally follow the red path all the way to v.

Next we deal with the case that all marked red edges are directed down. Now we ask what is the color of the highest marked edge, say e. If e is red, then we follow the entire blue path from f to x, then follow r^{-1} to b, then follow q to e and then follow the red path to v, where r is an evenlevel$(b;x)$ path.

Finally assume that e is blue. Again two cases arise: If e is directed down, then we follow the blue path from f to e, then follow q to the highest red edge (which is directed down), and then follow the red path to v. If e is directed up, then we follow the blue path from f to e, then follow q to b, then follow s to y and then follow the red path to v, where s is an evenlevel$(b;y)$ path. In all cases, a shorter even path to v is obtained, leading to a contradiction. □

Definition 24. (BFS honesty on p with respect to the base) Let p be an evenlevel(v) or oddlevel(v) path starting at unmatched vertex f to an arbitrary vertex v, let u lie on p with tenacity$(u) = t_m$ and let $b = \text{base}(u)$. Then we will say that u is *BFS honest on p w.r.t. b* if $p[b \text{ to } u]$ is an evenlevel$(b;u)$ (oddlevel$(b;u)$) path assuming $|p[b \text{ to } u]|$ is even (odd). Note that we are allowing b to appear either before or after u on path p.

Lemma 13 and Theorem 2 yield:

Lemma 14. *Let p be an* evenlevel(v) *path from unmatched vertex f to an arbitrary vertex v. Let $w \in p$ with* tenacity$(w) = t_m$ *and let $b = \text{base}(w)$. Then b also lies on p and w is BFS honest on p w.r.t. b.*

Proof : By Lemma 13 there is a vertex u on p such that $p[b \text{ to } u]$ contains all vertices in $p \cap (\mathcal{B}_{b,t_m} \cup \{b\})$ and $p[b \text{ to } u]$ is an evenlevel$(b;u)$ path. Let q be an evenlevel(b) path. Then by Corollary 4, $q \circ p[b \text{ to } u]$ is an evenlevel(u) path. Since w lies on this path and tenacity$(w) =$ tenacity$(u) = t_m$, by Theorem 2, w is BFS honest on this path. Hence w is BFS honest on p w.r.t. b. □

Note that in Lemmas 13 and 14, we took p to be an evenlevel(v) path; by Remark 2, this is without loss of generality. Observe that Lemma 13 allows b to appear either before or after u on path p. The following question arises: how does this affect whether b and u are BFS honest on p? Lemma 15 provides an answer and Example 16 illustrates the various cases.

[20] See Definition 18.

LEMMA 15. *Let p be an* evenlevel(v) *path from unmatched vertex f to an arbitrary vertex v such that there is a blossom $\mathcal{B}_{b,t_m}$ with $p \cap \mathcal{B}_{b,t_m} \neq \emptyset$. Let vertex $u \in (p \cap \mathcal{B}_{b,t_m})$ be such that $p[b$ to $u]$ contains all the vertices of $p \cap (\mathcal{B}_{b,t_m} \cup \{b\})$. Then the following hold:*

1. *If b appears before u on path p, then b and u are either both BFS honest or both not BFS honest on p.*
2. *If b appears after u on path p, then u is not BFS honest on p; furthermore, if b is BFS honest on p then $p[f$ to $b]$ is a* maxlevel(b) *path.*

Proof : **1).** The proof follows from Corollary 4.

2). If b appears after u on path p, then by Corollary 4, u is not BFS honest on p. Furthermore, if b is BFS honest on p then since $p[f$ to $b]$ cannot be a minlevel(b) = evenlevel(b) path, it must be an oddlevel(b) = maxlevel(b) path. □

EXAMPLE 16. In the graph of Figure 23, evenlevel$(u) = 8$; let p be this path. Vertex w appears on the p, tenacity$(w) = t_m = 13$, base$(w) = b$ and b appears before w on p. Observe that b and w are both BFS honest on p. Additionally, w is BFS honest on p w.r.t. b.

In the graph of Figures 25 and 26, tenacity$(u) = t_m = 11$ and base$(u) = b$. This graph has three bridges — of tenacity 11, 13 and 15. Observe that evenlevel$(v) = 16$, and let p denote the evenlevel(v) path; the three bridges appear in the order 15, 11, 13 on p. Now, b appears before u on p. Observe that b and u are both not BFS honest on p; however, u is BFS honest on p w.r.t. b.

In the graph of Figure 23, evenlevel$(c) = 12$; let q be the evenlevel(c) path. Observe that b appears after w on q and w is not BFS honest on q. However, b is BFS honest on q and $q[f$ to $b]$ is a maxlevel(b) path. Additionally, w is BFS honest on q w.r.t. b.

In the graph of Figure 23, evenlevel$(v) = 16$; let r be the evenlevel(v) path. Observe that b appears after w on r, and b and w are both not BFS honest on r; however, w is BFS honest on r w.r.t. b.

REMARK 4. As a result of the induction basis, we have established that every vertex v of tenacity t_m has a unique base, base(v). As stated in Section 8.1, the induction step, in Theorem 3, will enable us to establish an analogous fact about higher and higher tenacity vertices, eventually establishing it for every vertex of eligible tenacity. As a result, every such vertex w will have a unique base, base(w), which is an outer vertex and satisfies tenacity(base(w)) > tenacity(w). For ease of exposition, Definition 25 assumes that base(w) is well defined for each vertex w of eligible tenacity, in order to define the *iterated bases of* v, where tenacity$(v) = t_m$. A more "correct", though more cumbersome, way of doing this would be to define higher and higher iterated bases of v after each induction step.

DEFINITION 25. (Iterated bases of a vertex of tenacity t_m) For $v \in U_m$, let base$(v) = b$. Define the *first iterated base of* v to be b, denoted as follows: base$^1(v) = b$. When base(b), base(base(b)) etc. get defined in the induction step, we can define higher iterated bases of v. Thus, for $k \geq 1$, we will say that base$^{k+1}(v)$ = base(base$^k(v)$), assuming that base$^k(v)$ and base(base$^k(v)$) exist in the graph.

8.3. The Induction Step In this section, we will prove the induction step for Theorem 3; the basis of the induction was proved in Section 8.2.

Induction Hypothesis: Let t be an eligible tenacity, with $t_m + 2 \leq t < l_m$. Then each of the statements in Theorem 3 holds for tenacities in the range $[t_m, t-2]$.

After proving Statement 2, certain key notions will become well defined for the case of tenacity t. These definitions are formally stated after the proof of Statement 2 and they will be used for formally stating and proving the rest of the statements. These include the base of vertices of tenacity t, blossoms of tenacity t and iterated bases of a vertex.

In the induction step, we will need the following definition regarding the iterated bases of a vertex v such that $\text{tenacity}(v) < t$; these bases would already be defined in the previous iterations of the induction.

Definition 26. Let vertex v be such that $\text{tenacity}(v) < t$. Define $k(t,v)$ such that $\text{base}^{k(t,v)}(v)$ is the first iterated base of v having tenacity at least t, i.e.,

$$k(t,v) = \min\{l \mid \text{tenacity}(\text{base}^l(v)) \geq t\}.$$

Theorem 3. *Let t be an eligible tenacity, $U_t = \{u \in V \mid \text{tenacity}(u) = t\}$ and $v \in U_t$. The following hold:*

Statement 1: *Every* $\text{maxlevel}(v)$ *path contains a bridge of tenacity t.*

Statement 2: *The set $B(v)$ is a singleton.*

Statement 3: *Let $b = \text{base}(v)$. Then every* $\text{evenlevel}(v)$ $(\text{oddlevel}(v))$ *path consists of an* $\text{evenlevel}(b)$ *path followed by an* $\text{evenlevel}(b;v)$ $(\text{oddlevel}(b;v))$ *path, where the latter lies in $\mathcal{B}_{b,t} \cup \{b\}$*[21].

Statement 4: *The blossoms of tenacity t are disjoint and the set of blossoms of tenacity at most t forms a laminar family.*

Statement 5: *Let p be an* $\text{evenlevel}(u)$ *path from unmatched vertex f to an arbitrary vertex u. Let $w \in p$ with* $\text{tenacity}(w) = t$ *and let* $b = \text{base}(w)$. *Then b also lies on p and w is BFS honest on p w.r.t. b.*

Note that in Statement 5, we took p to be an $\text{evenlevel}(v)$ path; by Remark 2, this is without loss of generality.

Proof of Statement 1: The proof is given in Lemma 16 and is very different from the proof of the analogous fact, given in Lemma 5, in the induction basis. The reason is that the former needs to account for blossoms defined in the previous iterations of the induction.

Lemma 16. *Let v be a vertex of tenacity t and p be a* $\text{maxlevel}(v)$ *path. Then p contains a bridge of tenacity t.*

Proof : Assume that p starts at unmatched vertex f and let u be an arbitrary vertex on p. If $\text{tenacity}(u) = t$ then by Theorem 2, u is BFS honest on p and therefore $|p[f \text{ to } u]|$ is either $\text{minlevel}(u)$ or $\text{maxlevel}(u)$. If $\text{tenacity}(u) < t$ then by applying Lemma 20 from the previous induction step to p we get that $x = \text{base}^{k(t,u)}(u)$ lies on p and $p[x \text{ to } u]$ lies in $\{x\} \cup \mathcal{B}_{x,t-2}$, where $\mathcal{B}_{x,t-2}$ is the blossom of tenacity $t-2$ with base x.

Define sets S_1, S_2 and S as follows:

$$S_1 = \{w \mid w \text{ is on } p, \text{tenacity}(w) = t \text{ and } |p[f \text{ to } w]| = \text{maxlevel}(w)\}$$

$$S_2 = \{w \mid w \text{ is on } p, \text{tenacity}(w) < t \text{ and } \text{base}^{k(t,w)}(w) \text{ is higher than } w \text{ on } p\}$$

$$S = S_1 \cup S_2$$

For a definition of "higher/lower on a path" see Definition 18. Let w be the lowest vertex of S on p and let w' be the matched neighbor of w. First assume that $w \in S_2$; therefore $\text{tenacity}(w) < t$. Let $x = \text{base}^{k(t,w)}(w)$ and $w \in \mathcal{B}_{x,t-2}$. Applying Lemma 17 from the previous induction step we get that $w' \in \mathcal{B}_{x,t-2}$. Since w' is not the lowest vertex of S on p, we get that w must be odd w.r.t. p. Consider the following two cases.

Case 1: w is even with respect to p.
By the argument given above, $w \notin S_2$. Therefore $w \in S_1$ and hence $\text{tenacity}(w) = t$ and $|p[f \text{ to } w]| =$

[21] Note that $\mathcal{B}_{b,t}$, namely a blossom of tenacity t and base b, is defined in Definition 31, after proving Statement 2.

maxlevel(w) = evenlevel(w). By Lemma 1, tenacity(w') = tenacity$(w,w') = t$ and w' is lower than w on p. Since w' is not the lowest vertex of S on p, $w' \notin S$, and therefore $|p[f \text{ to } w']| = \text{minlevel}(w')$. Furthermore, since w' is odd w.r.t. p, w' is an inner vertex. Since maxlevel(w) = evenlevel(w), w is also an inner vertex. Therefore, the predecessors of w and w' are given by unmatched edges incident at them. Therefore neither is w a predecessor of w' nor is w' a predecessor of w. Hence (w,w') is a bridge of tenacity t.

Case 2: w is odd with respect to p.
Let (u,w) be the unmatched edge on p incident at w. Clearly u is lower than w on p and $|p[f \text{ to } u]|$ is even. We will show that (u,w) is a bridge of tenacity t.

First let us show that w is not a predecessor of u. Suppose tenacity$(u) \geq t$. By Theorem 2, u is BFS honest on p and since $u \notin S_1$, $|p[f \text{ to } u]| \neq \text{maxlevel}(u)$. Therefore $|p[f \text{ to } u]| = \text{minlevel}(u) = \text{evenlevel}(u)$. Furthermore the predecessor of u is its matched neighbor and not w.

Next suppose that tenacity$(u) < t$. Let $x = \text{base}^{k(t,u)}(u)$ and $u \in \mathcal{B}_{x,t-2}$. Clearly the predecessor of u is either x or it lies in the blossom $\mathcal{B}_{x,t-2}$. Applying Lemma 20 from the previous induction step to p we get that x lies on p and the path $p[x \text{ to } u]$ lies in $\{x\} \cup \mathcal{B}_{x,t-2}$. Since $u \notin S_2$, the base of $\mathcal{B}_{x,t-2}$, namely x, is lower than u on p. If $w \in S_1$, tenacity$(w) = t$ and therefore w does not lie in $\mathcal{B}_{x,t-2}$. If $w \in S_2$, the base of the blossom containing w is higher than w on p and again w does not lie in $\mathcal{B}_{x,t-2}$. Therefore, in both cases, w is not a predecessor of u. As stated above, by Lemma 20, $p[x \text{ to } u]$ lies in $\{x\} \cup \mathcal{B}_{x,t-2}$. Furthermore, since w does not lie in $\mathcal{B}_{x,t-2}$ and w is not the base of $\mathcal{B}_{x,t-2}$, we get that $p[x \text{ to } u]$ contains all vertices in $p \cap (\mathcal{B}_{x,t-2} \cup \{x\})$. Now by Lemma 19 of the previous induction step, $p[x \text{ to } u]$ is an evenlevel$(x;u)$ path[22]. Since tenacity$(x) \geq t$, by Theorem 2, x is BFS honest on p. Therefore $|p[f \text{ to } u]| = \text{evenlevel}(u)$ in this case as well.

Next, let us show that u is not a predecessor of w. If tenacity$(w) = t$ then $w \in S_1$ and $|p[f \text{ to } w]| = \text{maxlevel}(w)$. Therefore the predecessor of w is its matched neighbor and not u. If tenacity$(w) < t$ then $w \in S_2$ and $w \in \mathcal{B}_{y,t-2}$, where $y = \text{base}^{k(t,w)}(w)$. Therefore the predecessor of w lies in the blossom $\mathcal{B}_{y,t-2}$; again, u is not a predecessor of w. Hence (u,w) is a bridge.

Finally, we will show that tenacity$(u,w) = t$, thereby completing the proof. We will consider two cases. First assume that tenacity$(w) = t$. By Theorem 2, $|p[f \text{ to } w]| = \text{oddlevel}(w) = \text{evenlevel}(u) + 1$; the last equality follows from the fact that $|p[f \text{ to } u]| = \text{evenlevel}(u)$. Now,

$$\text{tenacity}(u,w) = \text{evenlevel}(u) + \text{evenlevel}(w) + 1 = \text{evenlevel}(u) + (t - \text{oddlevel}(w)) + 1 = t.$$

Next, assume that tenacity$(w) < t$. If so, $w \in S_2$ and $y = \text{base}^{k(t,w)}(w)$ is higher than w on p. Let q be a minlevel(v) path from f to v. Since tenacity$(v) = t$, $|p| + |q| = t$. Since tenacity$(y) \geq t$, by Theorem 2, y is BFS honest on p. Therefore $p[f \text{ to } y]$ is an oddlevel(y) path. Also, $q \circ p[v \text{ to } y]$ is an even path from f to y and therefore its length is at least evenlevel(y).

Now $|p[f \text{ to } y]| + |q \circ p[v \text{ to } y]| = |p| + |q| = t \geq \text{oddlevel}(y) + \text{evenlevel}(y) = \text{tenacity}(y) \geq t$, thereby implying that tenacity$(y) = t$ and $q \circ p[v \text{ to } y]$ is an evenlevel(y) path. Since $p[y \text{ to } w]$ contains all vertices in $p \cap (\mathcal{B}_{y,t-2} \cup \{y\})$, by Lemma 19 of the previous induction step, $p[y \text{ to } w]$ is an evenlevel$(y;w)$ path. Together with the previous assertion we get that $q \circ p[v \text{ to } w]$ is an evenlevel(w) path. Finally,

$$\text{tenacity}(u,w) = \text{evenlevel}(u) + \text{evenlevel}(w) + 1 = |p[f \text{ to } u]| + |q \circ p[v \text{ to } w]| + 1 = |p| + |q| = t.$$

□

Remark 5. As was done in Lemma 5, Lemma 16 can be strengthened to show uniqueness of the bridge as well. Since we will not need this fact, we will not prove it.

[22] Observe that the conclusion of Remark 7 does not apply here.

8.3.1. The Graphs H_t and H'_t For proving the induction step, we will define graphs H_t and H'_t, which are analogous to H_m and H'_m defined in Section 8.2.1. The main difference is that whereas all edges in the directed graph H_m are of unit length, those in H_t can be longer. The long edges help "jump" over lower tenacity blossoms.

Recall that $U_t = \{v \in V \mid \text{tenacity}(v) = t\}$. Let B_t denote the set of all bridges of tenacity t and let W_t denote the end points of all bridges in B_t.

Definition 27. For $v \in W_t$, define

$$v^* = \begin{cases} v & \text{if tenacity}(v) = t, \\ \text{base}^{k(t,v)}(v) & \text{if tenacity}(v) < t \end{cases}$$

Note that $k(t,v)$ is defined in Definition 26. We now explain the second case of the definition of v^*, i.e., if $\text{tenacity}(v) < t$. In this case, v is in a blossom of tenacity $t-2$ which must have been defined in the previous induction step. Now, v^* is meant to be the base of this blossom; it is given by $\text{base}^{k(t,v)}(v)$.

Let $W^*_t = \{v^* \mid v \in W_t\}$. For a vertex $v^* \in W^*_t$, let p denote an arbitrary $\text{minlevel}(v^*)$ path. Let V_t denote the set of all vertices of tenacity at least t on all $\text{minlevel}(v^*)$ paths p, for all vertices $v^* \in W^*_t$.

The next definition is related to Definition 20 and is motivated by Statement 5 of the induction hypothesis.

Definition 28. (The relations pred_t and pred^*_t) Let $v \in V_t$ with $\text{tenacity}(v) = t$ and let u be a predecessor of v; clearly, u may be unmatched. Define:

$$\text{pred}_t(v;u) = \begin{cases} u & \text{if tenacity}(u) \geq t, \\ \text{base}^{k(t,u)}(u) & \text{if tenacity}(u) < t \end{cases}$$

We will say that a vertex w *is* pred_t *of* v, denoted by $w = \text{pred}_t(v)$, if $\text{tenacity}(v) = t$ and there is a predecessor u of v such that $\text{pred}_t(v;u) = w$. The relation pred^*_t is recursively defined as follows: given vertices $u, v \in V_t$, with $\text{tenacity}(v) = t$, we will say that u *is* pred^*_t *of* v, denoted by $u = \text{pred}^*_t(v)$, if either $u = \text{pred}_t(v)$ or $u = \text{pred}^*_t(\text{pred}_t(v))$. Observe that if $u = \text{pred}^*_t(v)$ then $u \neq v$.

Next, we define directed graph H_t and undirected graph H'_t. Analogous to the definitions of graphs H_m and H'_m given in Section 8.2.1, the vertex sets of both H_t and H'_t are the same, namely V_t. The edge set of H_t, E_t, is defined as follows: For vertices $w, v \in V_t$, if $w = \text{pred}_t(v)$, then there is a directed edge $(v,w) \in E_t$. This edge is of unit length if $\text{pred}_t(v)$ is obtained from the first case of Definition 28 and it is longer if $\text{pred}_t(v)$ is obtained from the second case of Definition 28. In contrast, graph H_m, defined in Section 8.2.1, has unit length edges only. Define directed graph $H_t = (V_t, E_t)$.

Corresponding to the set of bridges of tenacity t, B_t, define

$$B^*_t = \{(u^*, v^*) \mid (u,v) \in B_t\}.$$

Define undirected graph $H'_t = (V_t, (B^*_t \cup E'_t))$, where the edge set E'_t is obtained from E_t by making each edge undirected. An edge $e \in E'_t$ is matched in H'_t if and only if the corresponding edge is present in G and is matched in G.

Example 17. Figures 19 and 20 illustrate the graphs H'_t corresponding to the graphs of Figures 14 and 26 for $t = 19$ and $t = 15$, respectively.

Remark 6. In Definition 27, for $v \in W_t$, if $\text{tenacity}(v) < t$, then $v^* = \text{base}^{k(t,v)}(v)$. This case adds an extra step in the process of finding an augmenting path, stated in Section 5.4.1; this step could not be described in that section due to lack of definition of v^*. We will explain this in the context of the graph of Figure 26. In order to find an $\text{evenlevel}(v)$ path, the algorithm needs to find

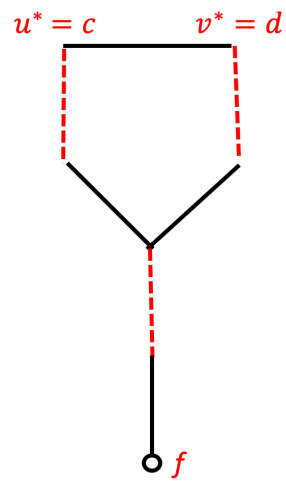

FIGURE 19. Graph H'_t for $t = 19$, corresponding to the graph of Figure 14.

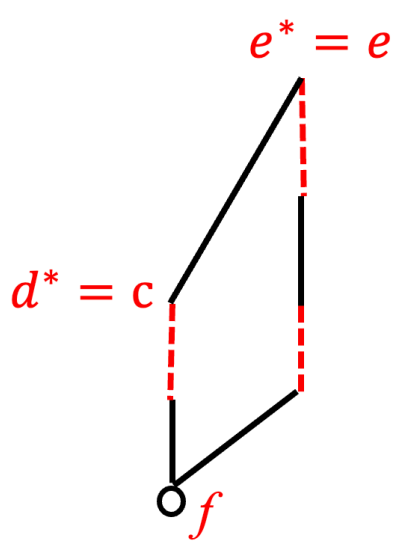

FIGURE 20. Graph H'_t for $t = 15$, corresponding to the graph of Figure 26.

a path from d to a in the blossom $\mathcal{B}_{f,15}$. Since $d^* \neq d$, it realizes that a part of this path, namely the path from d to $d^* = c$, needs to be found in the nested blossom $\mathcal{B}_{c,13}$.

The extra complexity that arises in the proof of the induction step is captured in the various definitions given above in this section. This complexity does not change the basic ideas needed for proving statements analogous to those given in Section 8.2 for the induction basis — the only difference in the formal statements is that subscript "m" is replaced by "t" throughout and t_m is

replaced by t. Below, we will summarize this development. The proof of Statement 2 is followed by some key definitions which are used by the rest of the statements of Theorem 3.

Lemmas 7 and 8 carry over and so does Procedure Bottleneck. The output of a run of this procedure on input $(u,v) \in B_t$ is a vertex b. As in the induction basis, this vertex is very special, and is called a *base*; a formal definition is given below. Clearly, b is an outer vertex. In general, H'_t will have a number of bases.

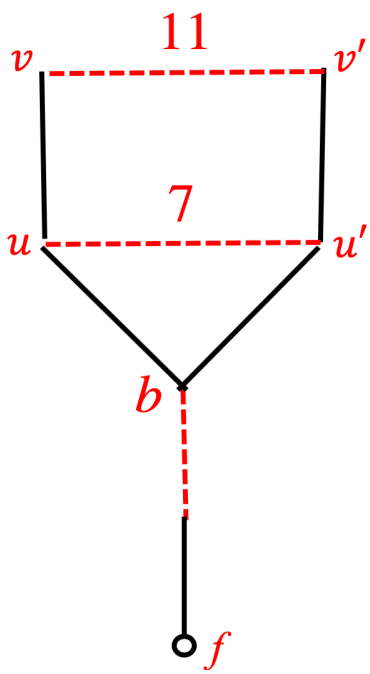

FIGURE 21. Vertex b is the base of u, u', v and v'. The tenacities of bridges (u,u') and (v,v') are indicated.

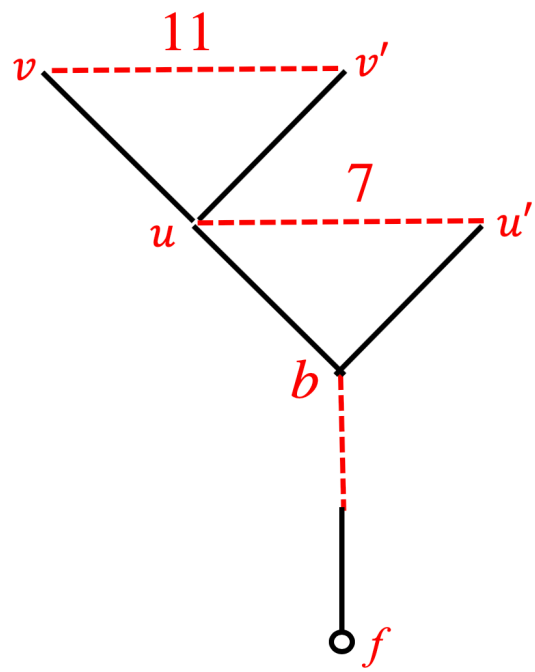

FIGURE 22. Vertex b is the base of u, u', v and v'. The tenacities of bridges (u,u') and (v,v') are indicated.

For each base b in H'_t, define the set

$$S_{b,t} = \{v \in U_t \mid b \text{ is pred}_t^* \text{ of } v\}.$$

Clearly $b \notin S_{b,t}$.

Proof of Statement 2: Lemma 10 also carries over and so does Corollary 2, thereby proving this statement.

Next, we formally define the base of vertices in U_t and blossoms of tenacity t.

DEFINITION 29. (The base of a vertex of tenacity t and basal vertices) For each $v \in U_t$, define base(v) to be the unique vertex, say b, in the set $B(v)$. We will say that the *base* of v is b. Each such vertex b will be called a *basal vertex*. Clearly, b is an outer vertex and tenacity$(b) > t$.

We can now define the iterated bases of a vertex of tenacity at most t; this is done in Definition 30. Once the entire induction step is proven, this definition holds for all vertices of eligible tenacity.

DEFINITION 30. (Iterated bases of a vertex of tenacity at most t) For $v \in U_t$, let base$(v) = b$. Define the *first iterated base of* v to be b, denoted as follows: base$^1(v) = b$. Next, consider a vertex u with tenacity$(u) < t$ and such that the induction hypothesis has established that base$^k(u) = v$, for $k \in \mathbb{Z}_+$. Then base$^{k+1}(u) = b$, i.e., base$^{k+1}(u) =$ base(base$^k(u)$).

EXAMPLE 18. In the graph of Figure 23, the iterated bases of vertex w are $\text{base}(w) = b$, $\text{base}^2(w) = b'$ and $\text{base}^3(w) = f$. Clearly, $\text{tenacity}(w) < \text{tenacity}(b) < \text{tenacity}(b') < \text{tenacity}(f)$.

We next come to the key definitions of blossom of tenacity t and the nesting of blossoms.

DEFINITION 31. (Blossom of tenacity t and base b) Let b be a basal vertex with $\text{tenacity}(b) > t$. Let $T_{b,t} = \{v \in U_t \mid \text{base}(v) = b\}$; observe that $T_{b,t} = S_{b,t}$. Then the *blossom of tenacity t and base b* is the set

$$\mathcal{B}_{b,t} = T_{b,t} \cup \left(\bigcup_{v \in (T_{b,t} \cup \{b\}),\ v \text{ is basal}} \mathcal{B}_{v,t-2} \right).$$

In the expression given above for $\mathcal{B}_{b,t}$, if for $v \in (T_{b,t} \cup \{b\})$, $\mathcal{B}_{v,t-2} \neq \emptyset$ then we will say that $\mathcal{B}_{v,t-2}$ is a *nested blossom of* $\mathcal{B}_{b,t}$. We will assume that the latter relation is transitively closed, i.e., if set A is a nested blossom of B and B is a nested blossom of C then A is a nested blossom of C. Clearly if $T_{b,t} = \emptyset$, then $\mathcal{B}_{b,t} = \mathcal{B}_{b,t-2}$.

The next fact is analogous to Corollary 3.

LEMMA 17. *Let (u,v) be a matched edge of tenacity t. Then u and v have the same base, say b, and both belong to the same blossom of tenacity t, namely $\mathcal{B}_{b,t}$.*

EXAMPLE 19. Figures 21 and 22 show two graphs with nested blossoms. Although the two "look different", in both graphs, vertex b is the base of vertices u, u', v and v'; the tenacities of bridges (u,u') and (v,v') are 7 and 11, respectively; the set $T_{b,11} = \{v, v'\}$; and the blossoms are $\mathcal{B}_{b,7} = \{u, u'\}$, $\mathcal{B}_{b,9} = \{u, u'\}$ and $\mathcal{B}_{b,11} = \{u, u', v, v'\}$.

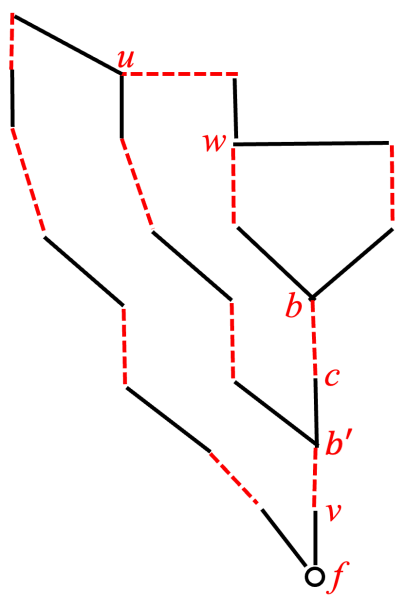

FIGURE 23. The iterated bases of vertex w are: $\text{base}(w) = b$, $\text{base}^2(w) = b'$ and $\text{base}^3(w) = f$, and those of vertex u are: $\text{base}(u) = b'$ and $\text{base}^2(u) = f$.

Definition 32. (Shortest path from a base to a vertex of tenacity t) Let $v \in U_t$ and $b = \text{base}(v)$. Then by an evenlevel$(b; v)$ (oddlevel$(b; v)$) path we mean a minimum even (odd) length alternating path in G from b to v that starts with an unmatched edge.

We will extend Definition 32 to define a *shortest path from an iterated base,* say b, to a vertex v of tenacity t. We will denote this path also by evenlevel$(b; v)$ (oddlevel$(b; v)$) depending on whether the path is even (odd) in length.

Proof of Statement 3: Claims analogous to Lemmas 10 and 11 hold, thereby proving this statement; it is analogous to Corollary 4 in the induction basis.

A claim analogous to Lemma 12 also holds:

Lemma 18. *Let $\mathcal{B}_{b,t}$ and $\mathcal{B}_{b',t}$ be two blossoms with bases $b \neq b'$. Then $\mathcal{B}_{b,t} \cap \mathcal{B}_{b',t} = \emptyset$.*

By Definition 31, if $\mathcal{B}_{b,t}$ and $\mathcal{B}_{b',t'}$ are two blossoms with $t > t'$, then either they are disjoint or the former contains the latter; note that we are allowing $b = b'$. This gives:

Corollary 5. *The set of blossoms of tenacity at most t forms a laminar family.*

Proof of Statement 4: Lemma 18 and Corollary 5 prove this statement.

Definition 33. (BFS honesty on p with respect to an iterated base) Let p be an evenlevel(v) path starting from unmatched vertex f to an arbitrary vertex v and let u lie on p with tenacity$(u) \leq t$. Let $b = \text{base}^k(u)$ be any one of the iterated bases of u that is well-defined at this stage of the induction. Then we will say that u is *BFS honest on p w.r.t. b* if $p[b \text{ to } u]$ is an evenlevel$(b; u)$ (oddlevel$(b; u)$) path assuming $|p[b \text{ to } u]|$ is even (odd). Note that we are allowing b to come either before or after u on p.

Proof of Statement 5: We will state a fact that is analogous to Lemma 13; its proof is also analogous and is omitted.

Lemma 19. *Let p be an* evenlevel(u) *path from unmatched vertex f to an arbitrary vertex u such that there is a blossom $\mathcal{B}_{b,t}$ with $p \cap \mathcal{B}_{b,t} \neq \emptyset$. Then the base of this blossom, b, also lies on p and there is a vertex $y \in (p \cap \mathcal{B}_{b,t})$ such that $p[b \text{ to } y]$ contains all vertices in $p \cap (\mathcal{B}_{b,t} \cup \{b\})$ and $p[b \text{ to } y]$ is an* evenlevel$(b; y)$ *path.*

One difference between Lemma 13 and Lemma 19 is that in the former, $b = \text{base}(u)$ whereas in the latter b may be any iterated base of y which is defined at this stage. Even so, Statement 5 follows from Lemma 19, along on the lines of Lemma 14. The reason is that the only iterated base of vertex w, occurring in Statement 5, which is defined at this stage is base(w), since tenacity$(w) = t$.

This completes the proof of Statement 5. Next let us integrate this statement over all the induction steps, until the current one, to get the following fact about iterated bases and nested blossoms. It will be used in the proof of Lemma 16 in the next induction step.

Lemma 20. *Let p be an* evenlevel(v) *path from unmatched vertex f to an arbitrary vertex v. Let vertex $u \in p$ with* tenacity$(u) \leq t$, *and let $l = k(t+2, u)$. Let the iterated bases of u,* $\text{base}^1(u), \ldots, \text{base}^l(u)$, *be $x_1, \ldots, x_l$, respectively. For $1 \leq i \leq l$, let $s_i = \text{tenacity}(x_i)$, and let $\mathcal{B}_{x_i, s_i - 2}$ be the blossom of tenacity $s_i - 2$ with base x_i. Then,*

1. *The iterated bases of u, $x_1, \ldots, x_l$, lie on p.*
2. *For $1 \leq i \leq l$, the path $p[x_i \text{ to } u]$ lies in $\{x_i\} \cup \mathcal{B}_{x_i, s_i - 2}$.*

Remark 7. In Lemma 20, u need not be BFS honest on p w.r.t. x_i, for $1 < i \leq l$; see Theorem 6 and Example 21.

Lemma 21, stated below, is analogous to Lemma 15. The various cases discussed in this lemma can be illustrated for any eligible tenacity $t > t_m$, similar to the way it was done in Example 16.

However, since the examples become large, we have illustrated only a subset of the cases in Example 20.

LEMMA 21. *Let p be an* evenlevel(v) *path from unmatched vertex f to an arbitrary vertex v such that there is a blossom $\mathcal{B}_{b,t}$ with $p \cap \mathcal{B}_{b,t} \neq \emptyset$. Let vertex $u \in (p \cap \mathcal{B}_{b,t})$ be such that $p[b$ to $u]$ contains all the vertices of $p \cap (\mathcal{B}_{b,t} \cup \{b\})$. Then the following hold:*

1. *If b appears before u on path p, then b and u are either both BFS honest or both not BFS honest on p.*
2. *If b appears after u on path p, then u is not BFS honest on p; furthermore, if b is BFS honest on p then $p[f$ to $b]$ is a* maxlevel(b) *path.*

This completes the proof of Theorem 3.

EXAMPLE 20. In the graphs of Figures 23 and 24, vertex $u \in \mathcal{B}_{b',15}$. Consider the evenlevel$(v)$ paths in both these graphs. On these paths, u appears before b'. Furthermore, u is not BFS honest and b' is BFS honest on these paths.

In the graph of Figure 24, consider the oddlevel(w) path. b' appears before u on this path and b' and u are both BFS honest on this path.

Finally, in the graph of Figure 26, consider the evenlevel(v) path, p. This path goes through the blossom $\mathcal{B}_{c,13}$; it enters the blossom at vertex d. d appears before c on this path and whereas d is not BFS honest, c is BFS honest on this path.

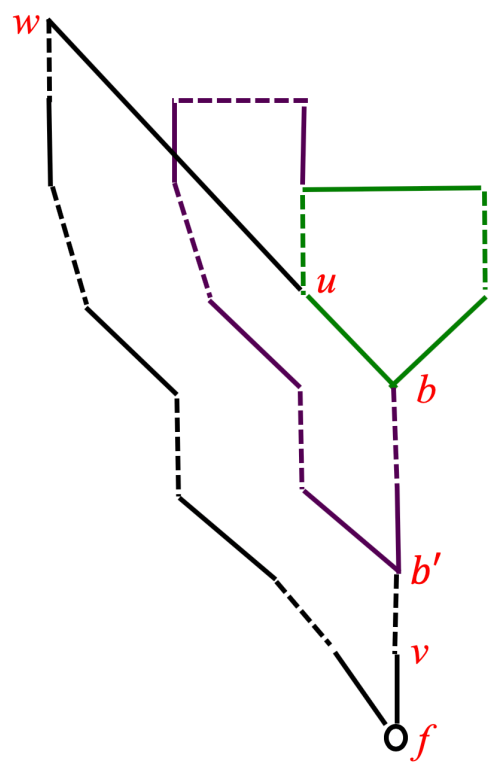

FIGURE 24. Let p be the evenlevel(v) path. Vertex u is BFS honest on p w.r.t. b and b'; however, u is not BFS honest on p w.r.t. f.

As observed in Example 14, vertices of tenacity l_m may have no base and as a result, Statements 2 to 5 of Theorem 3 do not hold for them. However, Statement 1 does hold and needs to be proven;

in particular, DDFS on the corresponding bridge will reveal an augmenting path if one exists. This case is singled out in the next theorem. Its proof is identical to that of Statement 1.

THEOREM 4. *For every vertex v of tenacity l_m, every* maxlevel(v) *path contains a bridge of tenacity l_m.*

8.4. BFS Honesty and Iterated Bases The properties established in Sections 8.2 and 8.3 help prove that the MV algorithm correctly finds the first minimum length augmenting path in a phase; in particular, the information left in the graph, such as pointers, is critical to accomplishing this task. After finding an augmenting path, the algorithm removes it and all vertices which cannot be present on a disjoint minimum length augmenting path.

This raises the following question: Does the remaining graph have the required structural properties and information to enable the algorithm to find successive augmenting paths? In this section, we prove additional properties which provide a positive answer to this question. These properties also show the sense in which minimum length alternating paths are not arbitrarily BFS dishonest, as stated in the Introduction.

DEFINITION 34. (The number of iterated bases of a vertex of eligible tenacity) Let v be a vertex of eligible tenacity. As noted earlier, the tenacities of its iterated bases keep increasing, i.e., $\text{tenacity}(\text{base}(v)) < \text{tenacity}(\text{base}^2(v)) < \text{tenacity}(\text{base}^3(v)) \ldots$. Let l be the smallest number such that $\text{tenacity}(\text{base}^l(v)) \geq l_m$; clearly, such a number exists, since the tenacity of unmatched vertices is at least l_m. Since $\text{base}^l(v)$ is not a vertex of eligible tenacity, by definition it does not have a base. We will say that *v has exactly l iterated bases.*

THEOREM 5. *Let v be a vertex of eligible tenacity, and let p be an* evenlevel(v) *or* oddlevel(v) *path, starting at unmatched vertex f, say. Assume that v has exactly l iterated bases. Then the following hold:*

1. *All l iterated bases of v lie on p; moreover, they occur in the order* $\text{base}^l(v), \text{base}^{l-1}(v), \ldots, \text{base}(v)$ *on p.*
2. *Each iterated base is BFS honest on p.*
3. *v is BFS honest on p w.r.t. each iterated base.*

Proof : We will prove by an induction on k, for $1 \leq k \leq l$, the following statement: $\text{base}^k(v)$ lies on p and $p[f \text{ to } \text{base}^k(v)]$ is an evenlevel$(\text{base}^k(v))$ path. This will establish the first two statements of the theorem and the third will then follow by Theorem 3.

Let $\text{base}(v) = b$. By Statement 3 of Theorem 3, every evenlevel(v) (oddlevel(v)) path consists of an evenlevel(b) path concatenated with an evenlevel$(b; v)$ (oddlevel$(b; v)$) path. Therefore $p[f \text{ to } b]$ is an evenlevel(b) path, hence establishing the basis of the induction.

Assume that the claim is true for k, where $1 \leq k < l$, and let $\text{base}^k(v) = u$. Then $p[f \text{ to } u]$ is an evenlevel(u) path. Let $\text{base}(u) = w$; clearly $\text{base}^{k+1}(v) = w$. Again by Statement 3 of Theorem 3, w lies on $p[f \text{ to } u]$ and $p[f \text{ to } w]$ is an evenlevel(w) path, hence establishing the induction step. □

THEOREM 6. *Let v be a vertex of eligible tenacity and p be an* evenlevel(v) *or* oddlevel(v) *path, starting at unmatched vertex f, say. Let u be a vertex of eligible tenacity which is on p and assume that u has exactly l iterated bases. Then the following hold:*

1. *If u is BFS honest on p then the l iterated bases of u satisfy the three conditions stated in Theorem 5.*
2. *If u is not BFS honest on p then:*
   (a) *All l iterated bases of u lie on p.*
   (b) *u is BFS honest on p w.r.t.* $\text{base}(u)$*, and for* $1 \leq k < l$, $\text{base}^k(u)$ *is BFS honest on p w.r.t.* $\text{base}^{k+1}(u)$.

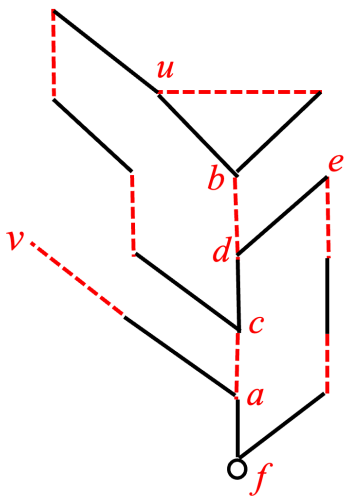

FIGURE 25. evenlevel(v) is finite; the reader is encouraged to find such a path.

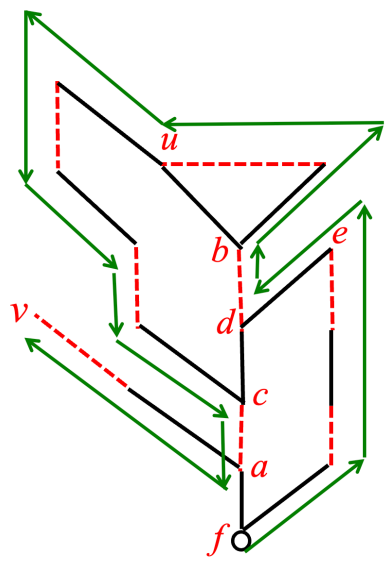

FIGURE 26. The evenlevel(v) path, p, is indicated. Vertices b and u are not BFS honest on p even though b occurs before u on p.

(c) base$^l(u)$ *is BFS honest on* p.

Proof :

1). Since $p[f$ to $u]$ is a minimum alternating path to u, the three conditions of Theorem 5 apply.

2). Next assume that u is not BFS honest on p. Even so, the first and second facts follow using Lemmas 13 and 19, and an easy induction on k. Since tenacity(base$^l(u)) \geq l_m$ and tenacity$(v) < l_m$, by Theorem 2, base$^l(u)$ is BFS honest on p, giving the third fact. □

EXAMPLE 21. In the graph of Figure 23, let p be the evenlevel(v) path; it contains vertices u and w which are both not BFS honest on p. The iterated bases of w are b, b' and f, and all three lie on p. Of these, b' and f are BFS honest on p, and b is not BFS honest on p; furthermore, each iterated base is BFS honest w.r.t. the next higher base, as proved in Theorem 6. Additionally, w is BFS honest w.r.t. b', w is not BFS honest w.r.t. f and b is not BFS honest w.r.t. f on p. The iterated bases of u are b' and f. Both lie on p and both are BFS honest on p, and b' is BFS honest on p w.r.t. f.

In the graph of Figure 26, the evenlevel(v) path, p, is indicated. The length of p is 16 and it contains vertex u. The iterated bases of u are b, c and f, and all three lie on p. Of these, c and f are BFS honest on p, and b is not BFS honest on p. Furthermore, each iterated base is BFS honest w.r.t. the next higher base; however, u is not BFS honest w.r.t. c, u is not BFS honest w.r.t. f and b is not BFS honest w.r.t. f on p.

9. Proof of Correctness and a Post-Mortem As stated in Section 1.1, one of the main ideas behind the MV algorithm is precise synchronization of events; this is described and proved

in Section 9.1. Section 9.3 proves that the MV algorithm correctly executes a phase and also establishes the running time of the algorithm. Finally, Section 9.4 provides a post-mortem by raising and answering the question, "Why is it essential to formalize such an elaborate purely graph-theoretic structure for proving correctness of the MV algorithm?"

9.1. Synchronization of Events

THEOREM 7. *Let t be an odd number with $t_m \le t \le l_m$, i.e., t is either an eligible tenacity or $t = l_m$. The following hold:*

1. *Algorithm 1 finds $Br(t)$, the set of bridges of tenacity t, by the end of execution of procedure MIN at search level i, where $t = 2i+1$.*
2. *For each vertex v such that* $\text{tenacity}(v) = t$, *Algorithm 1 assigns* $\text{minlevel}(v)$ *and* $\text{maxlevel}(v)$ *correctly.*

Proof : We will show, by strong induction on i, for $i = 0$ to $(l_m - 1)/2$, that at search level i, Algorithm 1 will accomplish:

Task 1: Procedure MIN assigns a minlevel of $i+1$ to exactly the set of vertices having this minlevel. It also identifies all props that assign a minlevel of $i+1$.

Task 2: By the end of execution of procedure MIN at this search level, $Br(2i+1)$ is the set of all bridges of tenacity $2i+1$.

Task 3: Procedure MAX assigns correct maxlevels to all vertices having tenacity $2i+1$.

This will establish both statements of the theorem.

The base case, $i = 0$, is obvious: MIN will assign an oddlevel of 1 to each neighbor of each unmatched vertex. Next we assume the induction hypothesis for all search levels less than i, and prove that Algorithm 1 will accomplish the three tasks at search level i.

Task 1: By the induction hypothesis, the minlevel assigned to vertex v at the beginning of execution of MIN at search level i is ∞ if and only if $\text{minlevel}(v) \ge i+1$. Since MIN searches from all vertices having level i along the correct parity edges and assigns a minlevel to a vertex only if its currently assigned minlevel is $\ge i+1$, any vertex v that is assigned a minlevel in this search level must indeed satisfy $\text{minlevel}(v) = i+1$, and the edge that reaches v will be correctly classified as a prop.

We next prove that every vertex v with $\text{minlevel}(v) = i+1$ will be assigned its minlevel in this search level, and every prop that assigns a minlevel of $i+1$ will be classified as a prop. Let $\text{minlevel}(v) = i+1$, let p be a $\text{minlevel}(v)$ path, and let (u,v) be the last edge on p. Clearly (u,v) is a prop, and every prop that assigns a minlevel of $i+1$ is of this type. Now, u must be BFS honest on p: If not, then v must occur on a shorter path to u, contradicting $\text{minlevel}(v) < i+1$. If $|p[f \text{ to } u]| = i = \text{maxlevel}(u)$ then $\text{tenacity}(u) < 2i+1$. Otherwise, $|p[f \text{ to } u]| = i = \text{minlevel}(u)$.

In either case, by the induction hypothesis, u has already been assigned a level of i. Therefore, at search level i, MIN will search from u along edge (u,v) and will find v. By the induction hypothesis, at this point, either the minlevel of v is set to either ∞ or $i+1$[23]. In either case, v will be assigned a minlevel of $i+1$, u will be declared a predecessor of v and (u,v) will be declared a prop.

Task 2: Let (u,v) be a matched bridge with $\text{tenacity}(u,v) = 2i+1$. By Lemma 1, $\text{tenacity}(u) = \text{tenacity}(v) = \text{tenacity}(u,v)$, and u and v are both inner. Therefore, $\text{oddlevel}(u) = \text{oddlevel}(v) = i$. Hence during search level i, MIN will determine that (u,v) is a bridge, that its tenacity is $2i+1$, and will insert it in $Br(2i+1)$.

Next assume that (u,v) is an unmatched bridge with $\text{tenacity}(u,v) = 2i+1$. By Lemma 3, if $\text{tenacity}(v) = \text{tenacity}(u,v) = 2i+1$, then v is an outer vertex and $\text{evenlevel}(v) \le i$. Therefore, the

[23] The latter case happens if $\text{evenlevel}(u) = i$ and v has been reached earlier in this search level while searching along an edge (u',v) with $\text{evenlevel}(u') = i$.

algorithm has already determined evenlevel(v). On the other hand, if tenacity(v) $< 2i+1$, then both its levels were determined by the end of the previous search level. By Lemma 3, these are the only two cases.

Therefore, in both cases, tenacity(u,v) will be ascertained by the end of execution of procedure MIN at search level i and $Br(2i+1)$ will be the set of all bridges of tenacity $2i+1$.

Task 3: Let tenacity(v) $= t$. By Statement 1 of Theorem 3 or by Theorem 4, depending on whether t is an eligible tenacity or $t = l_m$, v lies in the support of a bridge of tenacity $2i+1$, and by Task 2, this bridge is in $Br(2i+1)$ at the start of MAX in search level i. Therefore, DDFS will ascertain maxlevel(v) in this search level. □

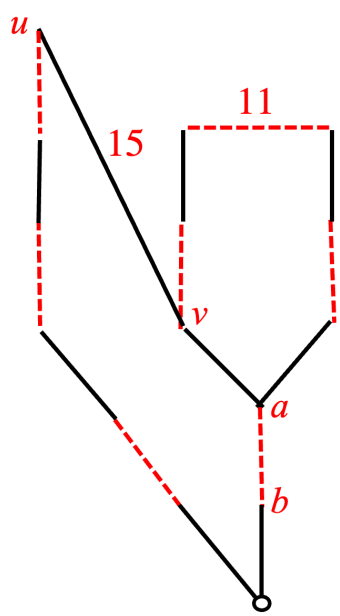

FIGURE 27. Can DDFS be performed on bridge (u,v) at search level 6?

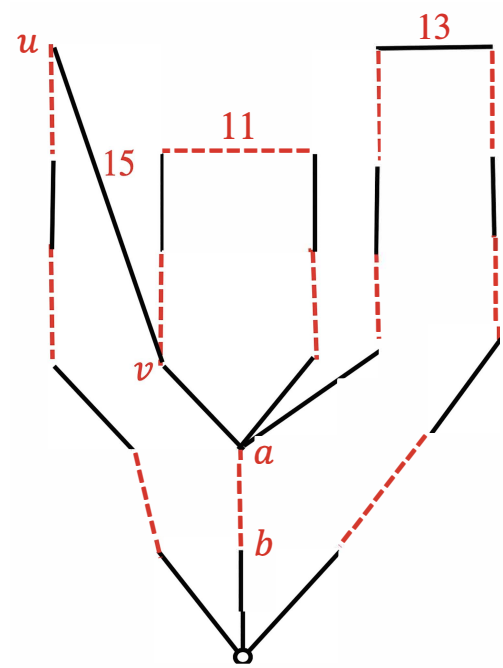

FIGURE 28. If so, vertices a and b will get wrong tenacities.

EXAMPLE 22. This example gives an insight into the idea of synchronizing of events. In Figure 27, the algorithm determines that (u,v) is a bridge of tenacity 15 at search level 6. However, according to Algorithm 1, DDFS should be performed on (u,v) at search level 7. The question arises, "Why wait till search level 7; why not perform DDFS on (u,v) when procedure MAX is run at search level 6?" In the graph of Figure 27 no mistakes will be made, provided the algorithm assigns tenacities of 15 vertices to a and b, i.e., the same as the tenacity of bridge (u,v).

However in the graph of Figure 28, when DDFS is performed on the bridge of tenacity 13 at search level 6, a and b will be assigned tenacities of 13; these are their correct tenacities in the graph of Figure 28. Next assume that bridge (u,v) is also processed at search level 6 and assume that it is processed before the bridge of tenacity 13; this is consistent with the arbitrary order in which bridges are processed. If so, a and b will be wrongly assigned tenacities of 15.

The synchronization of events imposed by the MV algorithm helps avoid such errors. Even though the tenacity of bridge (u,v) is determined at search level 6, DDFS on it is performed at search level 7, consistent with its tenacity. As a result, when the bridge of tenacity 13 is processed at search level 6, a and b are assigned their correct tenacities, i.e., 13. When processing bridge (u,v), DDFS will skip over the blossom of tenacity 13 and not encounter a and b.

Remark 8. The proper synchronization of events, described above, is central to guaranteeing that the MV algorithm can accomplish its task in linear time. In the absence of synchronization, if DDFS is executed for a bridge of tenacity t and it finds a vertex v, the only guarantee is that $\text{tenacity}(v) \leq t$, and determining its exact tenacity would be a time-consuming task. However, in the presence of synchronization, $\text{tenacity}(v) < t$ is not possible[24]. Therefore, in the presence of synchronization, if DDFS finds a vertex v, it is sure that $\text{tenacity}(v) = t$.

9.2. Relationship between Graph-Theoretic and Algorithmic Notions As stated in Section 5.2, the notions of petal and bud are intimately related to the notions of blossom and base; whereas the former are algorithmic notions, the latter are graph-theoretic. This relationship is formally established in Lemma 22. At the end of search level i, once MAX is done processing all bridges of tenacity $t = 2i + 1$, all blossoms of tenacity t can be identified via this lemma; its proof is straightforward and is omitted.

Lemma 22. *Let* $\text{tenacity}(v) = t$*, and at the end of search level* $i = (t-1)/2$*, assume that* $\text{bud}^*(v)$ *is* b*. Then* $\text{base}(v) = b$ *and the set* $T_{b,t}$ *defined in Definition 31 is:*

$$T_{b,t} = \{u \mid \text{tenacity}(u) = t \text{ and } \text{bud}^*(u) = b\}.$$

Furthermore, the blossom $\mathcal{B}_{b,t}$ *consists of the union of all petals whose* bud^* *is* b *at the end of search level* $i = (t-1)/2$*, together with each blossom of tenacity* $(t-2)$ *whose base is* b *or any of the vertices of these petals.*

Observe that if $\text{bud}^*(v)$ is computed at the end of search level $j > i$, then it may not be b anymore; however, it will be an iterated base of v.

9.3. Execution of a Phase and Proof of Running Time The proofs given in Section 9.1 show that the MV algorithm correctly finds one minimum length augmenting path in the given graph with an initial matching. Lemma 24 shows that it correctly finds a maximal set of such paths as well; it critically uses Lemma 23. Finally, Theorem 8 concludes with a proof of the running time.

Lemma 23. *For a blossom* $\mathcal{B}_{b,t}$*, if* b *is removed then procedure RECURSIVE REMOVE will remove all vertices of this blossom.*

Proof : We will prove the lemma by induction on the tenacity t of the blossom. The basis follows easily, since for each $v \in \mathcal{B}_{b,t_m}$, $b = \text{pred}^*_m(v)$, and therefore RECURSIVE REMOVE will remove the vertices of this blossom in order of increasing minlevel.

Next, assume that the statement is true for all blossoms of tenacity $t-2$, where $t_m < t < l_m$. Consider a blossom $\mathcal{B}_{b,t}$. For each vertex $v \in T_{b,t}$, $b = \text{pred}^*_t(v)$, see Definition 31. Moreover, the base of each blossom of tenacity $t-2$ nested in $\mathcal{B}_{b,t}$ is a vertex from $\{b\} \cup T_{b,t}$. Once b is removed, the vertices of $T_{b,t}$ will be removed in order of increasing minlevels, together with nested blossoms of tenacity $t-2$; the latter follows by the induction hypothesis. Hence $\mathcal{B}_{b,t}$ will be fully removed. □

Lemma 24. *The procedures given in Section 5.4 will find a maximal set of disjoint minimum length augmenting paths in* G*.*

Proof : Clearly, the first path, say p, found by the algorithm will be of length l_m, i.e., it will be a minimum length augmenting path. As argued earlier, the vertices removed by procedure

[24] If $\text{tenacity}(v) < t$, v would have already been assigned to a petal and the current DDFS would skip it.

RECURSIVE REMOVE of Section 5.4.2 cannot be part of a minimum length augmenting path that is disjoint from p.

The crux of the matter is: how do we guarantee that the remaining graph will "look like" a graph in which the first path is found, i.e., it has all the pointers and properties needed. The theorems of Section 8.4 guarantee that if a vertex $v \in p$ then each of its iterated bases is on p and will be removed. By Lemma 23, once the base of a blossom is removed, RECURSIVE REMOVE will remove all its vertices and therefore there will be no "half-eaten" blossoms in the graph when the next path needs to be found. The lemma follows. □

THEOREM 8. *The MV algorithm finds a maximum matching in general graphs in time $O(m\sqrt{n})$ on the RAM model and $O(m\sqrt{n} \cdot \alpha(m,n))$ on the pointer model, where α is the inverse Ackerman function.*

Proof : In a phase, each of the procedures MIN, MAX, finding augmenting paths, and RECURSIVE REMOVE examine each edge a constant number of times. The only other operation performed by the algorithm is that of computing bud* during DDFS. This can be implemented on the pointer model using the set union algorithm of Tarjan [31], which will take $O(m \cdot \alpha(m,n))$ time per phase. Alternatively, it can be implemented on the RAM model using the linear time algorithm for a special case of set union [10]; this will take $O(m)$ time per phase. Since $O(\sqrt{n})$ phases suffice for finding a maximum matching [15, 19], the theorem follows. □

REMARK 9. A question arising from Theorem 8 is whether there is a linear time implementation of bud* in the pointer model. [24] had claimed, without proof, that path compression by itself suffices to achieve this. They stated that because of the special structure of blossoms, a charging argument could be given that assigns a constant cost to each edge. This claim is left as an open problem. See a related open problem in Section 10.

Finally we note that since the MV algorithm consists of simple operations, involving no hidden constants, especially if implemented using the set union algorithm [31], it is not only fast in theory but also in practice. Over the years, several researchers have produced very fast implementations.

9.4. The role of graph-theoretic structural properties in the MV algorithm Finally we address the following question, "Why was it essential to formalize such an elaborate purely graph-theoretic structure for proving correctness of the MV algorithm?" Now that the reader is familiar with the structural definitions and claims, this question can be answered.

Assume that $\text{minlevel}(v) = i + 1$ and $\text{maxlevel}(v) = j + 1$, so that $\text{tenacity}(v) = i + j + 2 = t$, where t is an eligible tenacity or $t = l_m$. It is easy to see that there must be a neighbor, say u, of v, such that $\text{evenlevel}(u) = i$ or $\text{oddlevel}(u) = i$, depending on the parity of i. Therefore one step of breadth first search, while searching from u, will lead to assigning v its correct minlevel. We will say that u is the *agent that assigns v its minlevel*.

In contrast, none of the neighbors of v may have j as one of its levels. For instance, vertex b in the graph of Figure 5 has $\text{maxlevel}(b) = \text{oddlevel}(b) = 11$. Observe that $\text{evenlevel}(a) = \text{evenlevel}(c) = 8$ and $\text{evenlevel}(u) = 12$, i.e., none of the neighbors of b has an evenlevel of 10. The following questions arise:

1. What is the agent that assigns v its maxlevel and does it exist for each "relevant" vertex v?
2. Can this agent be found for each relevant vertex v?
3. How does this agent help assign v its maxlevel?

This paper provides very precise answers to all these questions. The agent that assigns v its maxlevel is a *bridge* whose tenacity equals $\text{tenacity}(v)$. Statement 1 of Theorem 3 and Theorem 4 prove that every $\text{maxlevel}(v)$ path contains such a bridge. Thus, in a sense, this was the central structural fact that needed to be proved.

"Relevant" vertices are those whose tenacity is eligible or is l_m. Theorem 3 proves that for each such vertex v, it is in the support of a bridge whose tenacity is tenacity(v), and Theorem 7 proves that such a bridge will be found "well in time". The answer to the third question is: execute the procedure of DDFS on the endpoints of this bridge.

These structural properties suffice for finding the first minimum length augmenting path. However, after its removal, could it be that the graph is left with "half-eaten blossoms" which simply do not support finding the next path, via the same process as the first one, even though a path exists? Lemma 24, which is based on additional properties established in Section 8.4, shows that subsequent paths don't use "half-eaten blossoms"; indeed, the procedure RECURSIVE REMOVE will remove all of them. Hence subsequent paths are found in the same way as the first one.

10. Discussion Matching is one of the "big three" problems in combinatorial optimization, along with linear programming and flow. Spectacular progress in this field has led to improved running times of the last two, as well as other fundamental problems, in the last three decades, e.g., see [29] as well as recent papers. Recent improvements in flow algorithms have also led to improved running times for the maximum matching problem in bipartite graphs: $O(m^{10/7})$ [23], $O(m^{4/3})$ [21] and $\tilde{O}(m+n^{1.5})$ [33]. Very recently, an almost linear time $O(m^{1+o(1)})$ algorithm was obtained [3]; however, it is currently unclear if this theoretical running time translates into very fast implementations for use in practice. A concerted effort has been made to improve the running time for general graph matching as well, but so far the MV algorithm has stood the test of time.

As is well-known, the general graph matching problem has numerous applications. We single out a particularly interesting and important one — to the kidney exchange matching market, which was first studied in [28]; see also the Scientific Background [4] for the 2012 Nobel Prize in Economics, awarded to Alvin Roth and Lloyd Shapley.

Assume that agent A requires a kidney transplant and agent B has agreed to donate one of her kidneys to A; however, their kidney types are not compatible. Assume further that (A', B') is another pair of people with an incompatibility. If it turns out that (A, B') and (A', B) are both compatible pairs, then let us say that the two pairs are *consistent*; if so, both transplants can be performed.

Next assume that a number of incompatible pairs are specified, $(A_1, B_1), \ldots, (A_n, B_n)$ and for every two pairs, we know whether they are consistent. Then the rudimentary problem[25] of finding the maximum number of disjoint consistent pairs reduces to maximum matching as follows. Let $G = (V, E)$ be a graph with $V = \{v_1, \ldots, v_n\}$ where v_i represents the pair (A_i, B_i), and $(v_i, v_j) \in E$ if and only if the two pairs (A_i, B_i) and (A_j, B_j) are consistent. Clearly a maximum matching in G will yield the answer.

In addition to the profound influence matching has had on the theory of algorithms (see the Introduction), it has also played a significant role in game theory and economics; we describe this next. The matching game forms one of the cornerstones of cooperative game theory, e.g., see [25], and its special case, the assignment game, forms a paradigmatic setting for studying the quintessential solution concept of the core[26] of a game [30].

Several matching-based problems form essential ingredients in the area of online and matching-based market design. Besides the application of general graph matching to kidney exchange (described above), two other variants of matching, namely stable matching [13] and online bipartite matching [18], lie at the core of this area, e.g., see [6]. The seminal 1962 paper of Gale and Shapley [13], on stable matching, initiated this area. With the advent of the Internet and mobile computing,

[25] Solution concepts proposed by economists are more involved, taking into account issues such as incentive compatibility, more general exchanges, etc., see [6].

[26] The recent paper [35] rectifies the fact that the core of the general graph matching game is empty via the notion of an approximate core.

it underwent a resurgence, leading to the launching of highly impactful new matching markets, e.g., the digital ads marketplaces, Uber, Lyft, Airbnb, Upwork and Match.com. The online bipartite matching problem has emerged as a paradigm for this area because of the online decision-making feature of these marketplaces.

Finally, here is an open question: Is there a linear time implementation of a phase of MV on the pointer model? Remark 9 states the approach mentioned in [24] for addressing this question. The following approach may be easier: Is it the case that the MV algorithm, implemented using the set union algorithm of Tarjan [31], actually runs in linear time per phase? The running time of $O(m\sqrt{n} \cdot \alpha(m,n))$, reported in Theorem 8, is based on imprecise assumptions that the number of calls to this data structure is $O(m)$ and number of elements manipulated is $O(n)$. More precise bounds on these two quantities are the number of calls to DDFS, i.e., the number of bridges of eligible tenacity, and the number of buds of petals[27] respectively. Using these facts and structural properties established in this paper, can one show that the total time devoted to set union in a phase is bounded by $O(m)$?

11. Acknowledgements I wish to thank Rohith Gangam and Ruta Mehta for diligently helping verify this proof, and Silvio Micali for embarking on a year-long journey, full of youthful exuberance, which led to the discovery of this algorithm. For the four-decade-long journey that led to the discovery of this proof, I must thank matching theory for its gloriously elegant combinatorial structure, which kept me going. Lastly, I wish thank the two referees for providing in-depth critiques of this paper, covering all aspects of the exposition and the proof.

References

[1] Ravindra K Ahuja, Thomas L Magnanti, and James B Orlin. *Network flows: theory, algorithms and applications*. Prentice Hall, 1995.

[2] Claude Berge. Two theorems in graph theory. *Proceedings of the National Academy of Sciences of the United States of America*, 43(9):842, 1957.

[3] Li Chen, Rasmus Kyng, Yang P Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. *arXiv preprint arXiv:2203.00671*, 2022.

[4] The Economic Sciences Prize Committee. Stable allocations and the practice of market design. *The Royal Swedish Academy of Sciences*, 2012. https://www.nobelprize.org/uploads/2018/06/advanced-economicsciences2012.pdf.

[5] Efim A. Dinitz. Algorithm for solution of a problem of maximum flow in networks with power estimation. In *Soviet Math. Doklady*, volume 11, pages 1277–1280, 1970.

[6] Federico Echenique, Nicole Immorlica, and Vijay V. Vazirani, editors. *Online and Matching-Based Market Design*. Cambridge University Press, 2023.

[7] J. Edmonds. Maximum matching and a polyhedron with 0,1-vertices. *Journal of Research of the National Bureau of Standards B*, 69B:125–130, 1965.

[8] Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17(3):449–467, 1965.

[9] Shimon Even and Oded Kariv. An $O(n^{2.5})$ algorithm for maximum matching in general graphs. In *16th Annual Symposium on Foundations of Computer Science*, pages 100–112, 1975.

[10] H. N. Gabow and R. E Tarjan. A linear-time algorithm for a special case of disjoint set union. *J. Comput. System Sci.*, 30:209–221, 1985.

[27] Observe that it suffices to maintain the set union structure over the set of buds of petals only: If during DDFS, a vertex v, in a previously constructed petal, is encountered, then the algorithm will follow pointers from v to its petal node to $\text{bud}(v)$. Next, the set union algorithm will find $\text{bud}^*(v)$. The work done to go from v to $\text{bud}(v)$ is charged to the edge that led to v and the rest of the work is charged to the set union algorithm.

[11] H. N. Gabow and R. E Tarjan. Faster scaling algorithms for general graph matching problems. *Journal of the ACM*, 38:815–853, 1991.

[12] Harold N Gabow. The weighted matching approach to maximum cardinality matching. *Fundamenta Informaticae*, 154(1-4):109–130, 2017.

[13] David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

[14] A. V. Goldberg and A. V. Karzanov. Maximum skew-symmetric flows and matchings. *Math. Program., Ser. A*, 100:537–568, 2004.

[15] J. Hopcroft and R. M. Karp. An $n^{5/2}$ algorithm for maximum matching in bipartite graphs. *SIAM Journal on Computing*, 2:225–231, 1973.

[16] M.R. Jerrum and A. Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18:1149–1178, 1989.

[17] M.R. Jerrum, L.G. Valiant, and V.V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986.

[18] Richard M Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 352–358, 1990.

[19] A. V. Karzanov. An exact estimate of an algorithm for finding a maximum flow, applied to the problem on representatives. *Problems in Cybernetics*, 5:66–70, 1973. Announced at the Seminar on Combinatorial Mathematics (Moscow, 1971).

[20] H.W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.

[21] Yang P Liu and Aaron Sidford. Faster energy maximization for faster maximum flow. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 803–814, 2020.

[22] L. Lovász and M.D. Plummer. *Matching Theory.* North-Holland, Amsterdam–New York, 1986.

[23] Aleksander Madry. Navigating central path with electrical flows: From flows to matchings, and back. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 253–262. IEEE, 2013.

[24] Silvio Micali and Vijay V. Vazirani. An $O(\sqrt{V}E)$ algorithm for finding maximum matching in general graphs. In *21st Annual Symposium on Foundations of Computer Science*, 1980.

[25] Hervé Moulin. *Cooperative microeconomics: a game-theoretic introduction*, volume 313. Princeton University Press, 2014.

[26] Marcin Mucha and Piotr Sankowski. Maximum matchings via Gaussian elimination. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 248–255. IEEE, 2004.

[27] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987.

[28] Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Pairwise kidney exchange. *Journal of Economic theory*, 125(2):151–188, 2005.

[29] A. Schrijver. *Theory of Linear and Integer Programming.* John Wiley & Sons, New York, NY, 1986.

[30] Lloyd S Shapley and Martin Shubik. The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130, 1971.

[31] R. E Tarjan. Efficiency of a good but not linear set union algorithm. *Journal of the ACM*, 22:215–225, 1975.

[32] L.G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201, 1979.

[33] Jan van den Brand, Yin-Tat Lee, Danupon Nanongkai, Richard Peng, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. Bipartite matching in nearly-linear time on moderately dense graphs. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 919–930. IEEE, 2020.

[34] Vijay V. Vazirani. A theory of alternating paths and blossoms for proving correctness of the $O(\sqrt{V}E)$ general graph maximum matching algorithm. *Combinatorica*, 14(1):71–109, 1994.

[35] Vijay V. Vazirani. The general graph matching game: Approximate core. *Games and Economic Behavior*, 132, 2022.

[36] Wikipedia. Isolation Lemma, . https://en.wikipedia.org/wiki/Isolation_lemma.

[37] Wikipedia. Harold W. Kuhn, . https://en.wikipedia.org/wiki/Harold_W._Kuhn.