

# A Nearly Linear-Time Distributed Algorithm for Exact Maximum Matching

Taisuke Izumi\*      Naoki Kitamura\*      Yutaro Yamaguchi\*

## Abstract

In this paper, we propose a randomized  $\tilde{O}(\mu(G))$ -round algorithm for the maximum cardinality matching problem in the CONGEST model, where  $\mu(G)$  means the maximum size of a matching of the input graph  $G$ . The proposed algorithm substantially improves the current best worst-case running time. The key technical ingredient is a new randomized algorithm of finding an augmenting path of length  $\ell$  with high probability within  $\tilde{O}(\ell)$  rounds, which positively settles an open problem left in the prior work by Ahmadi and Kuhn [DISC'20].

The idea of our augmenting path algorithm is based on a recent result by Kitamura and Izumi [IEICE Trans.'22], which efficiently identifies a sparse substructure of the input graph containing an augmenting path, following a new concept called *alternating base trees*. Their algorithm, however, resorts to a centralized approach of collecting the entire information of the substructure into a single vertex for constructing an augmenting path. The technical highlight of this paper is to provide a fully-decentralized counterpart of such a centralized method. To develop the algorithm, we prove several new structural properties of alternating base trees, which are of independent interest.

## 1 Introduction

### 1.1 Background and Our Result

In this paper, we consider the maximum matching problem<sup>1</sup> for general graphs in the CONGEST model, a standard computational model for distributed graph algorithms. The network is modeled as an undirected graph  $G = (V, E)$  of  $n$  vertices (nodes) and  $m$  edges (communication links). Each vertex executes a deployed algorithm following round-based synchrony, and each edge can transfer a message of  $O(\log n)$  bits per round bidirectionally. In distributed settings, any graph problem admits the trivial solution of aggregating all the topological information of  $G$  into a single vertex (and solving the problem by any centralized algorithm). Due to the bandwidth limitation of the CONGEST model, such a trivial solution requires  $\Theta(m)$  rounds, which may be  $\Theta(n^2)$ . Hence, the technical challenge in the CONGEST model is to design an algorithm such that each node outputs the solution without the whole input information at hand.

The maximum matching problem for general graphs is often referred to as the “hardest” problem of those admitting a polynomial-time solution in the context of centralized graph algorithms. A similar situation can be seen in the CONGEST model: Until recently, it was unclear if one could construct an algorithm faster than the trivial  $O(n^2)$ -round algorithm or not. Looking at the lower-bound side, it has been proved that any CONGEST algorithm of computing exact maximum matchings requires  $\Omega(\sqrt{n} + D)$  rounds in the worst case even if the input is restricted to

\*Osaka University. Email: izumi.taisuke.ist@osaka-u.ac.jp, {n-kitamura, yutaro.yamaguchi}@ist.osaka-u.ac.jp

<sup>1</sup>Throughout this paper, we only consider the unweighted problem, and thus simply refer to a maximum cardinality matching as a *maximum* matching.

bipartite graphs [2], where  $D$  is the diameter of the input graph. One may think that this bound can be improved when considering general graphs. However, it is also shown that such a lower bound cannot be improved even for general graphs as long as the proof is based on the two-party communication complexity argument, which is a dominant technique to prove lower bounds in the CONGEST model [4]. In 2022, Kitamura and Izumi [19] proposed an algorithm running within  $\tilde{O}(\mu(G)^{1.5})$  rounds, where  $\mu(G)$  represents the maximum matching size of the input graph. It is the first algorithm beating the worst-case bound of  $O(n^2)$  rounds. However, the current gap between the best upper and lower bounds is still large.

The primary contribution of this paper is the development of a new CONGEST algorithm for the maximum matching problem running in nearly linear time. The main theorem is stated below:

**Theorem 1.** *There exists a randomized CONGEST algorithm to compute a maximum matching of a given graph with high probability<sup>2</sup> that terminates within  $\tilde{O}(\mu(G))$  rounds.*

## 1.2 Technical Highlight

Our algorithm is built on the top of several known techniques. We first present an outline of those techniques, and then explain the key technical idea newly proposed in this paper.

**Edmonds’ Blossom Algorithm** The starting point is the well-known centralized algorithm by Edmonds, the so-called *blossom algorithm* [11]. This algorithm iteratively finds *augmenting paths* and improves the current matching size using the identified augmenting paths. If an augmenting path is found, the current matching is augmented by flipping the labels of matching and non-matching edges along the path. The seminal theorem by Berge [7] states that the current matching is the maximum if and only if there exists no augmenting path. Hence, the central difficulty of the maximum matching problem lies in the efficient detection of augmenting paths.

Whereas a simple approach of expanding alternating paths in the BFS-like manner correctly works in bipartite graphs, it fails in general graphs because of the existence of odd cycles, which destroys the very useful property holding in bipartite graphs that any subpath of a shortest alternating path is also a shortest alternating path. Following [25, 26], we refer to this property as the *BFS-honesty*. Edmonds’ algorithm basically works in the same way as the case of bipartite graphs, i.e., starting from unmatched vertices, it makes a BFS-like forest grow up by choosing matching and non-matching edges alternately. If the forest growth process encounters a *blossom*, an odd cycle and an alternating path from the tree source to the vertex incident to two non-matching edges in the cycle, the algorithm contracts its odd-cycle part into a single vertex; subsequently, the forest growth process continues in the resultant graph. An important observation is that the existence of augmenting paths is preserved by this contraction operation. Once an augmenting path of the contracted graph is found, an augmenting path of the original graph must be recovered from it. Roughly, when the augmenting path finally found contains a vertex associated with a contracted cycle, that vertex is expanded to an appropriate subpath in the cycle. This process is not so simple because contracted cycles can be nested (i.e., some odd cycle may contain a vertex associated with a contracted cycle), but certainly possible.

Since Edmonds’ algorithm is essentially sequential, it is highly challenging to obtain its efficient distributed implementation in the CONGEST model. There are three major obstacles in the implementation. 1) The iterative improvement of matchings: It can be iterated  $\Theta(n)$  times. 2) The construction of BFS-like forests: The depth of a tree in the forest may become  $\Theta(n)$ . 3) Graph

---

<sup>2</sup>Throughout this paper, the phrase “with high probability” means it succeeds with probability at least  $1 - 1/n^c$  for an arbitrary constant  $c > 1$ .

contraction: The simulation of the contracted graph on the original message-passing topology can incur expensive time cost, and the path recovery process is inherently sequential, particularly when contracted cycles are nested. As a result, even it is nontrivial to compute an augmenting path with a running time linearly dependent on its length. In fact, a natural distributed implementation of Edmonds’ approach incurs  $\Theta(\ell^2)$  rounds for finding an augmenting path of length  $\ell$ . Since the length of a shortest augmenting path can become  $\Theta(n)$ , such an implementation is not sufficient to beat the trivial bound of  $O(n^2)$  rounds.

**Framework Based on Hopcroft–Karp Analysis** In addition to faster augmenting-path detection, efficient distributed implementation of Edmonds’ algorithm includes yet another challenge that the iterative improvement can be repeated  $\Theta(n)$  times. To design a nearly linear-time algorithm, we must attain the amortized  $\tilde{O}(1)$ -round complexity of finding one augmenting path, which looks a big hurdle. Fortunately, it is achieved by employing the classic Hopcroft–Karp structural property on the length of shortest augmenting paths:

**Theorem 2** (Hopcroft and Karp [16]). *Let  $G = (V, E)$  be an undirected graph, and  $\mu(G)$  be the maximum matching size of  $G$ . Given a matching  $M \subseteq E$  of  $G$  with  $|M| \leq \mu(G) - k$ , there always exists an augmenting path of length less than  $\lfloor 2\mu(G)/k \rfloor$ .*

Following this theorem, it suffices to find an augmenting path of length  $\ell$  within  $\tilde{O}(\ell n^{1-\epsilon})$  rounds for improving the trivial  $O(n^2)$ -round upper bound. By an elementary analysis, such an algorithm works as an amortized  $\tilde{O}(n^{1-\epsilon})$ -round algorithm. In particular, the algorithm with  $\epsilon = 1$  (i.e.  $\tilde{O}(\ell)$ -round algorithm) is the best-possible goal along this approach, which attains the amortized optimal  $\tilde{O}(1)$ -round complexity.

**Distributed Maximum Matching Verification by Ahmadi and Kuhn** Motivated by the aforementioned framework, Ahmadi and Kuhn [1] provided an interesting idea toward  $\tilde{O}(\ell)$ -round augmenting path algorithms. They actually did not present any construction algorithm, but a *verification* algorithm of determining if a given matching is maximum or not. To be more precise, the proposed algorithm provides each vertex with information on the lengths of shortest odd- and even-alternating paths from a closest unmatched vertex, where an odd-alternating (resp. even-alternating) path is an alternating path with an odd (resp. even) length. The running time of the algorithm is parameterized by  $\ell$ , and all the lengths till  $\ell$  are reported within  $O(\ell)$  rounds.

The computed information can be used to identify a subgraph which contains exactly two unmatched vertices admitting a shortest augmenting path connecting them. It implies that the quest for  $\tilde{O}(\ell)$ -round augmenting path construction algorithms can be considered in a much more moderate setting, where the two endpoints of the target shortest augmenting path has been known, and every vertex knows the lengths of shortest odd- and even-alternating paths from the endpoints. However, it should be emphasized that constructing an augmenting path is still challenging even in this setting. Due to a lack of the BFS-honesty, subpaths of a shortest-alternating path are not necessarily shortest alternating paths. Then, the straightforward path-tracing approach based on distance information does not work.

**Semi-Distributed Augmenting Path Construction by Kitamura and Izumi** Kitamura and Izumi [19] addressed the limitation of the aforementioned research line. They proposed an algorithm which finds an augmenting path within  $\tilde{O}(\ell\sqrt{n})$  rounds on the top of the Ahmadi–Kuhn framework, if the current matching admits an augmenting path of length  $\ell$ . Their approach is

to construct a *sparse certificate*, which is a sparse (i.e., containing  $O(n)$  edges<sup>3</sup>) subgraph of  $G$  preserving the reachability from unmatched vertices by alternating paths. Specifically, a sparse certificate contains a  $\theta$ -alternating path from an unmatched vertex  $s$  to some vertex  $t$  for  $\theta \in \{\text{odd}, \text{even}\}$  if and only if the original graph admits such a path. A node can collect all the information on the sparse certificate within  $O(n)$  rounds, trivially allowing the centralized solution of finding augmenting paths. The  $\tilde{O}(\ell\sqrt{n})$  running time is attained by using this algorithm if  $\ell \geq \sqrt{n}$ , or the naive  $O(\ell^2)$ -round algorithm otherwise.

It is easily seen that such a certificate always exists: Looking at Edmonds' algorithm, the augmenting path finally found in the contracted graph is expanded to a path in the original graph whose edges lie in the constructed BFS-like forest and all the contracted odd cycles. In other words, the union of the BFS-like forest and all the contracted cycles, whose number is at most  $n$ , becomes a sparse certificate. The key ingredient of the Kitamura–Izumi algorithm is an efficient identification of such a sparse certificate in the setting of the Ahmadi–Kuhn framework. It first constructs a rooted spanning tree  $T$  called an *alternating base tree*, which corresponds to the BFS-like forest by Edmonds, within  $O(1)$  rounds using the distance information provided by the Ahmadi–Kuhn framework (it should be noted that  $T$  is a single tree because every vertex already knows the endpoints of the target augmenting path with the aid of the Ahmadi–Kuhn framework, and thus it suffices to start the tree growth process from one of the endpoints). To transform the alternating base tree into a sparse certificate, they introduced the notion of *edge levels* and showed that the desired structure is obtained by adding one minimum-level outgoing non-tree edge for each subtree. It is also shown that the computation of edge levels and the identification of minimum-level outgoing non-tree edges are efficiently implemented.

**Our Technique** While our algorithm is based on the techniques by Kitamura and Izumi, it includes many new ideas and insights regarding the problem. A big leap from Kitamura–Izumi is that we provide a *fully-decentralized* path-construction algorithm. All the approaches explained above leave one big obstacle faced in decentralized path construction: how is the task of finding long augmenting paths decomposed into that of finding shorter subpaths? While such a mechanism seems crucial for accelerating the construction of augmenting paths, nothing other than the sequential tree growth process from one endpoint has been known so far, at least in the CONGEST model. In addition, the (advanced) use of the Kitamura–Izumi algorithm also faces another obstacle: The augmenting path preserved in the constructed sparse certificate is not necessarily the shortest. Hence, it is also unlikely that one can benefit from their sparsification technique, except for the centralized collection approach.

The primary technical contribution of this paper is to demonstrate that both obstacles can be removed just by a small modification of their sparsification algorithm. More precisely, we use a slightly modified definition of edge levels, and consider the same sparse certificate construction. Under our new definition, it can be shown that the constructed certificate preserves at least one *shortest*  $\theta$ -alternating path for any relevant pair of vertices and any  $\theta \in \{\text{odd}, \text{even}\}$  (if it exists), including a shortest augmenting path. Furthermore, the preserved alternating paths are equipped with an important property: let  $T$  be the alternating base tree,  $s$  be a vertex, and  $t$  be a descendant of  $s$  in  $T$  and reachable from  $s$  by a  $\theta$ -alternating path. Then it is guaranteed that at least one preserved shortest  $\theta$ -alternating path  $P$  from  $s$  to  $t$  is split into the prefix  $P_1$  and the suffix  $P_2$  respectively lying at the outside and the inside of the subtree of  $T$  rooted by  $t$ . Intuitively, this property implies that the task of finding a shortest  $\theta$ -alternating path  $P$  from  $s$  to  $t$  is decomposed

---

<sup>3</sup>The original paper [19] actually constructs the certificate of  $O(\mu(G))$  edges, and attains the algorithm running within  $\tilde{O}(\ell\sqrt{\mu(G)})$  rounds. Here we explain the idea with the simplification of  $\mu(G) = \Theta(n)$ .

into two independent subtasks of finding  $P_1$  and  $P_2$  at the outside/inside of the subtree. Following this observation, we develop a surprisingly simple recursive algorithm of constructing a shortest augmenting path, which runs in time linearly dependent on the output length.

To show the correctness our algorithm, many rich structural properties of alternating base trees are newly proved in this paper. We believe that these insights are of independent interest, and provide a very useful toolbox toward further improvement of algorithms, which will run in sublinear time, and other related settings such as approximate and/or weighted variants.

### 1.3 Related Works

There have been many known results for the maximum matching problem in divergent settings [3, 5, 9, 15, 17, 18, 21, 22]. Table 1 summarizes the known (approximate) maximum matching algorithms in distributed settings, with focus on relevant to our result. Many literatures have considered approximate maximum matchings in both CONGEST and the more relaxed LOCAL model, wherein no restriction of bandwidth is incurred. The main interest of those studies is the localization of the computation process. Typically, the results along this line exhibits the trade-off between the approximation factor and the running time, and some algorithms attain the  $(1 - \epsilon)$ -approximation factor, which can be seen as an exact algorithm by taking  $\epsilon < 1/n$ . However, most of those results in the CONGEST model have the running time exponentially depending on  $1/\epsilon$ , and thus their performance as an exact algorithm is very poor. One exception is the recent breakthrough result by Fischer, Slobodan, and Uitto [12], which provides a  $\text{poly}(1/\epsilon, \log n)$ -time  $(1 - \epsilon)$ -approximate algorithm. Recently, a similar result for weighted matchings was proposed by Huang and Su [17]. However, the hidden exponents of  $1/\epsilon$  for those algorithms are quite larger than two, and thus it is not sufficient to break the trivial  $O(n^2)$  bound. The paper by Huang and Su [17] also provides an extensive survey on distributed approximate maximum matching algorithms.

Focusing on the exact computation, there are three algorithms of attaining nontrivial upper bounds. Ben-Basat, Kawarabayashi, and Schwartzman [6] presents a CONGEST algorithm whose running time is parameterized by  $\mu(G)$ . The running time of the algorithm is  $\tilde{O}(\mu(G)^2)$ . They also proved the lower bound of  $\Omega(\mu(G))$  in the LOCAL model. A generalized  $\Omega(1/\epsilon)$  lower bound for the  $(1 - \epsilon)$ -approximate maximum matching is also presented. While these bounds apply to the CONGEST model, the worst-case instance of providing these bounds has the diameter  $\Theta(n)$ . Parameterizing this lower bound with diameter  $D$ , it only exhibits the bound of  $\Omega(D)$ . Since the  $\tilde{\Omega}(\sqrt{n} + D)$ -round bound we mentioned applies to the hard-core instances of  $D = O(\log n)$ , their bounds are incomparable. An improved  $\tilde{O}(n)$ -round upper bound for bipartite graphs was proposed by Ahmadi, Kuhn, and Oshman [2]. This is also the first paper utilizing the framework based on the analysis of Hopcroft–Karp. It is worth remarking that no  $o(n)$ -round algorithm is known even for bipartite matchings. The third algorithm is the one by Kitamura and Izumi [19] attaining  $\tilde{O}(\mu(G)^{1.5})$  rounds. Until our result, this was only the known result faster than the trivial  $O(n^2)$ -round algorithm.

We also remark centralized exact maximum matching algorithms. Edmonds' blossom algorithm is the first centralized polynomial-time solution for the maximum matching problem [10, 11]. Hopcroft and Karp [16] proposed a phase-based algorithm of finding multiple augmenting paths. Their algorithm finds a maximal set of pairwise disjoint shortest augmenting paths in each phase. They showed that  $O(\sqrt{n})$  phases suffice to obtain a maximum matching and proposed an  $O(m)$ -time per phase implementation for bipartite graphs. Several studies have reported phase-based algorithms for general graphs that attain  $O(\sqrt{nm})$  time [8, 13, 24–26]. While our algorithm actually finds multiple disjoint augmenting paths in parallel, it cannot benefit from such a parallelism because finding the last augmenting path can take  $\Omega(n)$  rounds after all. It is an interesting open

Table 1: Lower and upper bounds for (semantically) exact maximum matching algorithms in the CONGEST model

	Time Complexity	Approximation Level	Remark
Ben-Basat et al. [6]	$\Omega(\mu(G))$	exact	LOCAL
Ben-Basat et al. [6]	$\Omega\left(\frac{1}{\epsilon}\right)$	$1 - \epsilon$	LOCAL
Ben-Basat et al. [6]	$\tilde{O}(\mu(G)^2)$	exact	
Kitamura and Izumi [19]	$\tilde{O}(\mu(G)^{1.5})$	exact	
Ahmadi et al. [2]	$\tilde{O}(\mu(G))$	exact	bipartite
Fischer et al. [12]	$O(\text{poly}(\epsilon, \log n))$	$1 - \epsilon$	
Huang and Su [17]	$O(\text{poly}(\epsilon, \log n))$	$1 - \epsilon$	weighted
Ahmadi et al. [2]	$O\left(\frac{\log^2 \Delta + \log^* n}{\epsilon^2}\right)$	$1 - \epsilon$	bipartite
Ahmadi et al. [2]	$\hat{\Omega}(\sqrt{n} + D)$	exact	bipartite
<b>Our result</b>	$\tilde{O}(\mu(G))$	<b>exact</b>	

question if the technique of the phase-based algorithms above can be utilized for obtaining much faster CONGEST algorithms or not.

## 2 Preliminaries

### 2.1 Notations and Terminologies

We denote the vertex set and edge set of a given graph  $G$  by  $V(G)$  and  $E(G)$ , respectively. The diameter of a graph  $G$  is denoted by  $D(G)$ . We denote by  $I_G(v)$  the set of edges incident to  $v$  in  $G$ . Throughout this paper, we only consider simple undirected graphs, and an edge connecting two vertices  $u$  and  $v$  is denoted by  $\{u, v\}$ .

A *path* of  $G$  is an alternating sequence  $P = v_0, e_1, v_1, e_2, \dots, e_\ell, v_\ell$  of vertices and edges such that all the vertices are distinct and  $e_i = \{v_{i-1}, v_i\} \in E(G)$  for every  $1 \leq i \leq \ell$ . A path is often regarded as a subgraph of  $G$ . For a path  $P = v_0, e_1, v_1, \dots, v_\ell$  of  $G$  and an edge  $e = \{v_\ell, u\} \in E(G)$  satisfying  $u \notin V(P)$ , we denote by  $P \circ e$  or  $P \circ u$  the path obtained by adding  $e$  and  $u$  to the tail of  $P$ . Similarly, for another path  $P'$  starting at  $v_\ell$  with  $V(P) \cap V(P') = \{v_\ell\}$ , we denote by  $P \circ P'$  the path obtained by concatenating  $P'$  to the tail of  $P$  (without duplication of  $v_\ell$ ). The inversion of the path  $P = v_0, e_1, v_1, \dots, v_\ell$  (i.e., the path  $v_\ell, e_\ell, v_{\ell-1}, e_{\ell-1}, \dots, v_0$ ) is denoted by  $\bar{P}$ . The length  $\ell$  of a path  $P$  is denoted by  $|P|$ . Given a path  $P$  containing two vertices  $u$  and  $v$  such that  $u$  precedes  $v$ , we denote by  $P[u, v]$  the subpath of  $P$  starting from  $u$  and terminating with  $v$ .

For a graph  $G$ , a *matching*  $M \subseteq E(G)$  is a set of edges that do not share endpoints. We denote by  $\mu(G)$  the maximum size of a matching of a graph  $G$ . In the following definitions, we fix a matching  $M$  of  $G$ . A vertex  $v$  is called *matched* if  $M$  intersects  $I_G(v)$ , and *unmatched* otherwise. A path  $P = v_0, e_1, v_1, e_2, \dots, e_\ell, v_\ell$  is called *alternating* if exactly one of  $e_{i-1}$  and  $e_i$  is in  $M$  for every  $1 < i \leq \ell$ . An *augmenting path* is an alternating path connecting two different unmatched vertices. We say that  $(G, M)$  has an augmenting path if there exists an augmenting path in  $G$  with respect to  $M$ . For notations  $V(G)$ ,  $E(G)$ ,  $D(G)$ , and  $\mu(G)$ , we often use the shorthand without argument  $G$  unless it causes any ambiguity.

Given a subgraph  $H$  of  $G$  and a matching  $M$  of  $H$ , an alternating path with an odd (resp. even)

length is called an *odd-alternating path* (resp. *even-alternating path*). For  $u, v \in V(H)$ ,  $v$  is called *odd-reachable* (resp. *even-reachable*) from  $u$  in  $H$  if there exists an odd-alternating (resp. even-alternating) path from  $u$  to  $v$ . A vertex  $v$  is simply called *reachable* from  $u$  if it is odd-reachable or even-reachable from  $u$ . We denote the length of the shortest odd (resp. even)-alternating path from  $u$  to  $v$  in  $H$  by  $\text{dist}_H^{\text{odd}}(u, v)$  (resp.  $\text{dist}_H^{\text{even}}(u, v)$ ). If  $v$  is not odd (resp. even)-reachable from  $u$ , we define  $\text{dist}_H^{\text{odd}}(u, v) = \infty$  (resp.  $\text{dist}_H^{\text{even}}(u, v) = \infty$ ). For  $\theta \in \{\text{odd}, \text{even}\}$ ,  $\bar{\theta}$  represents the parity different from  $\theta$ . The parity  $\theta \in \{\text{odd}, \text{even}\}$  satisfying  $\text{dist}_H^\theta(u, v) < \text{dist}_H^{\bar{\theta}}(u, v)$  (i.e., the parity of the shortest alternating path from  $u$  to  $v$ ) is denoted by  $\gamma_H(u, v)$ . Notice that it is well-defined for any reachable  $v$  because  $\text{dist}_H^\theta(u, v) \neq \text{dist}_H^{\bar{\theta}}(u, v)$  necessarily holds as they have different parities. We also introduce the notion of parity  $\rho(e)$  of each edge  $e \in E_G$ , which is defined as  $\rho(e) = \text{odd}$  if  $e$  is a matching edge, and  $\rho(e) = \text{even}$  otherwise. Intuitively,  $\rho(e)$  represents the parity of the alternating path  $P$  from  $f$  such that  $P \circ e$  can become an alternating path.

While a majority of notations and terminologies defined above actually depend on the matching  $M$  currently given,  $M$  is fixed in most of the following arguments, and thus we introduce them without describing an argument  $M$  for simplicity. In addition, the subscripts  $H$  of notations  $\text{dist}_H^\theta(u, v)$  and  $\gamma_H(u, v)$  are also abbreviated if  $H = G$ .

## 2.2 CONGEST Model

This paper assumes the CONGEST model, which is one of the standard models in designing distributed graph algorithms. A distributed system is represented by a simple connected undirected graph  $G$  of  $n$  vertices and  $m$  edges. Vertices and edges are uniquely identified by  $O(\log n)$ -bit integer values, and the computation follows round-based synchronous message passing. In one round, each vertex  $v$  sends and receives  $O(\log n)$ -bit messages through each of its incident edges, and executes local computation following its internal state, local random bits, and all the received messages. A vertex can send different messages to different neighbors in one round, and all the messages sent in a round are necessarily delivered to their destinations within the same round.

Throughout this paper, we assume that each vertex initially knows the value of  $n$  and a 2-approximate upper bound  $\hat{\mu}$  of  $\mu(G)$ . This assumption is not essential because both values are easily computed by an  $O(\mu(G))$ -round preprocessing presented in [19].

## 3 $O(\ell)$ -Round Construction of Augmenting Path

### 3.1 Specifications of Frameworks

As stated in the introduction, the technical core of our algorithm is to develop an  $O(\ell)$ -round algorithm of constructing an augmenting path of length at most  $\ell$  on the top of the Ahmadi–Kuhn framework. Such a construction is transformed into an  $\tilde{O}(\mu(G))$ -round maximum matching algorithm in the CONGEST model following the framework by the Hopcroft–Karp analysis. We first present the formal specification of these two frameworks. On the Hopcroft–Karp framework, we have the following lemma:

**Lemma 1** (Ahmadi, Kuhn, and Oshman [2]). *Assume a CONGEST algorithm  $A(M, \ell)$  with the following property:*

*For any graph  $G = (V, E)$  and matching  $M \subseteq E$ ,  $A(M, \ell)$  finds a nonempty set of vertex-disjoint augmenting paths with high probability if  $(G, M)$  has an augmenting path of length at most  $\ell$ , and it always terminates within  $O(\ell)$  rounds. Each vertex  $u \in V(G)$*

outputs the predecessor and successor of the output augmenting path to which  $u$  belongs (if it exists).

Then there exists a randomized CONGEST algorithm computing a maximum matching within  $\tilde{O}(\mu(G))$  rounds in total.

The framework by Ahmadi and Kuhn is formalized by the following lemma:

**Lemma 2** (Ahmadi and Kuhn [1], Kitamura and Izumi [19]). *Assume that there exists an algorithm  $A'(M, \ell)$  satisfying the specification of Lemma 1 under the following additional assumptions.*

- *The graph  $G$  contains exactly two unmatched vertices  $f$  and  $g$ , which satisfies  $\text{dist}^{\text{odd}}(f, g) \leq \ell$ .*
- *Every vertex  $v \in V(G)$  has a shortest alternating path from  $f$  of length at most  $\ell$ , i.e.,  $v$  satisfies  $\text{dist}^{\gamma(f,v)}(f, v) \leq \ell$ . In addition,  $v$  knows  $\text{dist}^{\text{odd}}(f, v)$ ,  $\text{dist}^{\text{even}}(f, v)$ , and the ID of  $f$ . This assumption implies the diameter of  $G$  is  $O(\ell)$ .*

Then there exists a randomized algorithm  $A(M, \ell)$  satisfying the specification of Lemma 1 without these assumptions.

While the proofs are given in the cited literatures, they do not state the lemmas explicitly in the forms as above. We provide a complementary explanation of the proofs in Appendix A. In virtue of these lemmas, we can focus only on the construction of  $A'$  of Lemma 2. Note that the algorithm  $A'$  we present is deterministic, i.e., random bits are used only in the part of Ahmadi–Kuhn. Hence, we do not pay much attention to the success probability of the algorithm.

In the following argument, we assume that the input graph  $G$  satisfies the assumption of Lemma 2, and treat  $A'$  as an algorithm of finding an augmenting path from  $f$ . The notations  $\text{dist}^\theta(u, v)$  and  $\gamma(u, v)$  are abbreviated to  $\text{dist}^\theta(v)$  and  $\gamma(v)$  if  $u = f$ .

### 3.2 Alternating Base Tree and Edge Level

Our algorithm uses the concept of alternating base trees proposed by Kitamura and Izumi [19].

**Definition 1** (Alternating base tree). *An alternating base tree for  $G$ ,  $M$ , and  $f$  is a rooted spanning tree  $T$  of  $G$  satisfying the following conditions:*

- *$f$  is the root of  $T$ .*
- *For any  $v \in V(G) \setminus \{f\}$  and its parent  $u$ ,  $\text{dist}^{\gamma(v)}(v) = \text{dist}^{\overline{\gamma(v)}}(u) + 1$  holds.*

It is easy to check that such a spanning tree always exists. As a vertex might have two or more shortest alternating paths,  $T$  is not uniquely determined (see Figure 1 (1) and (2) for examples). In the following argument, we fix one alternating base tree  $T$  arbitrarily chosen. The alternating base tree does not necessarily contain an alternating path from  $f$  to each vertex  $v$ . For example, both alternating base trees in Figure 1 have no alternating path from  $f$  to  $v_9$ .

We denote by  $\text{p}(v)$  the parent of  $v$  in  $T$ , and denote by  $T_v$  the subtree of  $T$  rooted by  $v$ . We define the *level* of each edge as follows:

**Definition 2** (Edge level). *For any  $e = \{u, v\} \in E(G) \setminus E(T)$ , we define the sum-level and max-level of  $e$  as  $\text{level}_{\text{sum}}(e) = \text{dist}^{\rho(e)}(u) + \text{dist}^{\rho(e)}(v)$  and  $\text{level}_{\text{max}}(e) = \max(\text{dist}^{\rho(e)}(u), \text{dist}^{\rho(e)}(v))$ , respectively. The level of an edge  $e \in E(G)$ , denoted by  $\text{level}(e)$ , is defined as the ordered pair  $(\text{level}_{\text{sum}}(e), \text{level}_{\text{max}}(e))$  if  $e \in E(G) \setminus E(T)$ , and as  $(0, 0)$  if  $e \in E(T)$ .*



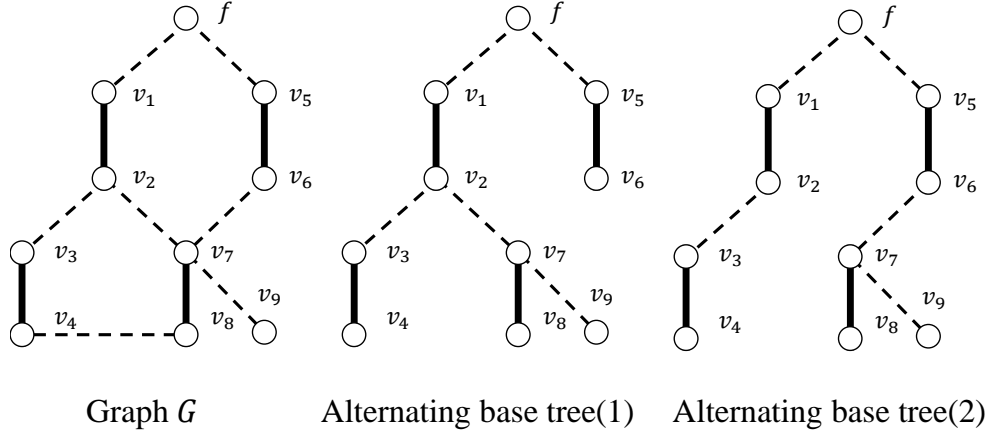


Figure 1: Examples of alternating base trees (from [19]). Solid lines are matching edges, and broken lines are non-matching edges.

We introduce the comparison operator  $<$  of edge levels following their lexicographic order, i.e., for any two edges  $e_1, e_2 \in E(G)$ ,  $\text{level}(e_1) < \text{level}(e_2)$  holds if and only if either  $\text{level}_{\text{sum}}(e_1) < \text{level}_{\text{sum}}(e_2)$  or  $(\text{level}_{\text{sum}}(e_1) = \text{level}_{\text{sum}}(e_2)) \wedge (\text{level}_{\text{max}}(e_1) < \text{level}_{\text{max}}(e_2))$  holds. Given an alternating path  $P$ , we refer to the maximum level of the edges in  $P$  (with respect to  $<$ ) as the *level* of  $P$ , which is denoted by  $\text{level}(P)$ . The *sum-level* of  $P$  and its notation  $\text{level}_{\text{sum}}(P)$  are also defined similarly.

The notion of edge levels is illustrated in Figure 2. A finite level of an edge  $e$  implies that there exists an odd alternating walk starting and ending at  $f$  and containing  $e$ . Suppose that one of such walks forms a blossom, i.e., an odd cycle containing  $e$  and an alternating path connecting  $f$  and the odd cycle, called the *stem*, that share exactly one vertex. Then the sum-level of  $e$  corresponds to the “minimum volume” of such a blossom, where the volume is measured by the number of edges in the blossom with double count of those in the stem. The max-level of  $e$  is an upper bound for the “height” of a minimum-volume blossom  $B$ , which is defined as the distance from  $f$  to the farthest vertex along  $B$ . Especially, the minimum max-level of edges in  $B$  is equal to the height. If we take a minimum-volume blossom  $B$  among all the blossoms containing  $e$  in the cycle, for any vertex  $w$  in the cycle except for the one shared with the stem, the two paths from  $f$  to  $w$  along  $B$  are the shortest odd- and even-alternating paths to  $w$  respectively.

**Definition 3** (Minimum outgoing edge (MOE) of  $T_v$ ). *Let  $v \in V(G)$  be a vertex. An edge which is not contained in  $T$  (not restricted to  $T_v$ ) and crosses  $V(T_v)$  and  $V(G) \setminus V(T_v)$  is called an outgoing edge of  $T_v$ . A minimum outgoing edge (MOE) of  $T_v$  is an outgoing edge of  $T_v$  whose level is minimum.*

For simplicity of the proof, for each vertex  $v$  such that  $T_v$  has no outgoing edge (including the case of  $v = f$ ), we introduce a virtual outgoing edge  $e_v$  of  $T_v$  whose level is  $(\infty, \infty)$ . While a minimum outgoing edge is not necessarily unique, we choose one of them for each  $v \in V(G)$  as the *canonical MOE* of  $T_v$  following the rules below:

- If  $T_v$  has no outgoing edge, the virtual edge  $e_v$  is chosen.
- If there exist MOEs incident to  $v$ , an arbitrary one of them is chosen.
- Otherwise, an arbitrary MOE is chosen.




Figure 2: An example of edge levels. Two values paired by parentheses respectively mean the values of  $\text{dist}^{\text{odd}}(\cdot)$  and  $\text{dist}^{\text{even}}(\cdot)$ . The values paired by brackets represent edge levels. All tree edges have level  $[\infty, \infty]$ , which is not explicitly described in the figure.

We denote the canonical MOE of  $T_v$  by  $\text{out}(v)$ , and  $E^*(T_v)$  as the set of all outgoing edges of  $T_v$  with a finite edge level. In the following argument, we denote the two endpoints of  $\text{out}(v)$  by  $y_v$  and  $z_v$ , i.e.,  $\text{out}(v) = \{y_v, z_v\}$ . Using this notation, we always assume  $z_v \in V(T_v)$ . We also use the shorthand notation  $\rho(v) = \rho(\text{out}(v))$ .

The key technical observation is the following lemma (the proof is deferred to Section 4).

**Lemma 3.** *For any  $t \in V(G)$  such that  $E^*(T_t) \neq \emptyset$ , the following three statements hold (see Fig. 3):*

- (S1) *There exists a shortest  $\rho(t)$ -alternating path  $Y'$  from  $f$  to  $y_t$  of level less than  $\text{level}(\text{out}(t))$ .*
- (S2) *There exists a shortest  $\rho(t)$ -alternating path  $Z'$  from  $f$  to  $z_t$  of level less than  $\text{level}(\text{out}(t))$ .*
- (S3) *The path  $X' = Y' \circ \{y_t, z_t\} \circ \overline{Z'}[z_t, t]$  is a shortest  $\overline{\gamma(t)}$ -alternating path from  $f$  to  $t$ . This does not depend on the choices of  $Y'$  and  $Z'$  in the above two statements.*

An illustrative explanation of this lemma is provided in Figure 3. Any bireachable vertex  $t$  is contained in the odd cycle of some blossom  $B$  “strictly” (i.e.,  $t$  is not the vertex shared with the stem of  $B$ ). This cycle must contain an edge in  $E^*(T_v)$  because the parent of  $t$  is the predecessor of the shortest  $\gamma(t)$ -alternating path to  $t$ , and thus the  $\gamma(t)$ -alternating path along  $B$  must cross between  $V(T_v)$  and  $V(G) \setminus V(T_v)$  at least once. Lemma 3 implies that at least one blossom containing  $\text{out}(t)$  and  $t$  crosses exactly once. In addition, since  $\text{out}(t)$  has the minimum sum-level of all edges in  $E^*(T_v)$ , some blossom containing  $\text{out}(t)$  is minimum of all ones containing  $t$  (in their cycles strictly). Hence the  $\gamma(t)$ - and  $\overline{\gamma(t)}$ -alternating paths along such a blossom is shortest.

### 3.3 Algorithm Details

We present our algorithm in the centralized manner, and discuss its distributed implementation in Section 3.4. We first introduce the notion of *extendability* of alternating paths.

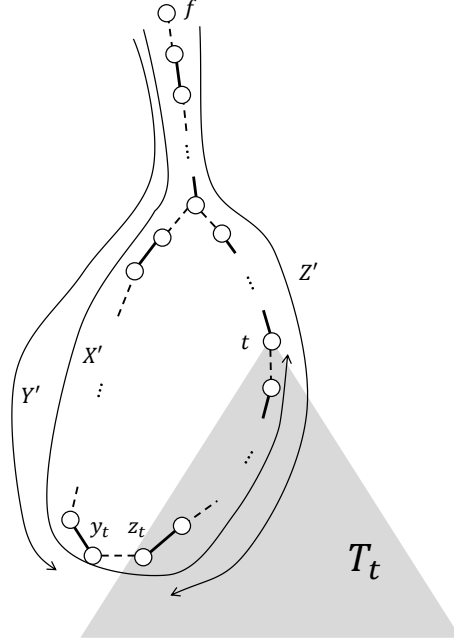


Figure 3: A sketch of Lemma 3.

**Definition 4.** For  $s, t \in V(G)$  such that  $s$  is an ancestor of  $t$ , and  $\theta \in \{\text{odd}, \text{even}\}$ , a pair  $(s, t)$  is called  $\theta$ -extendable if there exists a shortest  $\theta$ -alternating path  $P$  from  $f$  to  $t$  satisfying  $s \in V(P)$  and  $\text{level}(P) < \text{level}(\text{out}(s))$ . Then the path  $P$  satisfying the condition above is called a  $\theta$ -extension of  $(s, t)$ , and  $P[s, t]$  is called a  $\theta$ -extendable path from  $s$  to  $t$ .

A key feature of extendable paths is that the task of finding minimum-level extendable paths is self-reducible. The procedure  $\text{EXTPATH}(s, t, \theta)$  recursively computes a minimum-level  $\theta$ -extendable path from  $s$  to  $t$  following such a reduction scheme. The pseudocode of this procedure is presented in Algorithm 1. The principle behind the algorithm is explained as follows:

- For  $\theta = \gamma(t)$ : One can show that  $(s, \mathbf{p}(t))$  is  $\bar{\theta}$ -extendable, and hence the algorithm recursively invokes  $\text{EXTPATH}(s, \mathbf{p}(t), \bar{\theta})$ , which returns a minimum-level  $\bar{\theta}$ -extendable path  $Y$  from  $s$  to  $\mathbf{p}(t)$ . It is easy to show that  $Y$  does not contain  $t$ . Thus, the algorithm finally returns  $Y \circ \{\mathbf{p}(t), t\}$ , which is shown to be a minimum-level  $\theta$ -extendable path from  $s$  to  $t$ .
- For  $\theta = \overline{\gamma(t)}$ : Since any  $\overline{\gamma(t)}$ -alternating path from  $f$  to  $t$  cannot contain the edge  $\{\mathbf{p}(t), t\}$ , any  $\theta$ -extension of  $(s, t)$  must contain an outgoing edge of  $T_t$  to reach  $t$ . That is,  $\text{level}(\text{out}(t)) < \text{level}(\text{out}(s))$  holds (recall that any  $\theta$ -extension of  $(s, t)$  has a level less than  $\text{level}(\text{out}(s))$ ). This means that both  $y_t$  and  $z_t$  are in  $T_s$ . By Lemma 3, there exist two alternating paths  $Y' \circ \{y_t, z_t\}$  and  $Z'$  from  $f$  to  $z_t$  whose levels are at most  $\text{level}(\text{out}(t)) < \text{level}(\text{out}(s))$ . Since  $\{\mathbf{p}(s), s\}$  is the only edge of level less than  $\text{level}(\text{out}(s))$  to reach the vertices in  $V(T_s)$ , both  $Y'$  and  $Z'$  necessarily contain  $\{\mathbf{p}(s), s\}$ . Hence,  $Y'$  and  $Z'$  are respectively  $\rho(t)$ -extensions of  $(s, y_t)$  and  $(t, z_t)$ , i.e., they are  $\rho(t)$ -extendable. The algorithm recursively invokes  $\text{EXTPATH}(s, y_t, \rho(t))$  and  $\text{EXTPATH}(t, z_t, \rho(t))$  for identifying two minimum-level  $\rho(t)$ -extendable paths,  $Y$  from  $s$  to  $y_t$  and  $Z$  from  $t$  to  $z_t$ . Due to the existence of  $Y'[s, y_t]$  and  $Z'[t, z_t]$ , whose levels are less than  $\text{level}(\text{out}(t))$ , the levels of  $Y$  and  $Z$  are also less than  $\text{level}(\text{out}(t))$ . Then  $Y$  lies at the outside of  $T_t$ , and  $Z$  lies at the inside of  $T_t$ . It implies that the concatenated path

$Y \circ \{y_t, z_t\} \circ \bar{Z}$ , which is the output of the algorithm, is an alternating path from  $s$  to  $t$ . Utilizing the statement (S3) of Lemma 3, one can prove that the output is a minimum-level  $\theta$ -extendable path.

---

**Algorithm 1** Procedure  $\text{EXTPATH}(s, t, \theta)$

---

```

1: if  $s = t$  then
2:   return the path of length zero
3: else if  $\theta = \gamma(t)$  then
4:    $Y \leftarrow \text{EXTPATH}(s, \mathbf{p}(t), \bar{\theta})$ 
5:   return  $Y \circ \{\mathbf{p}(t), t\}$ 
6: else
7:    $Y \leftarrow \text{EXTPATH}(s, y_t, \rho(t))$  //  $\text{out}(t) = \{y_t, z_t\}$ , where  $z_t$  is in  $T_t$ 
8:    $Z \leftarrow \text{EXTPATH}(t, z_t, \rho(t))$ 
9:   return  $Y \circ \{y_t, z_t\} \circ \bar{Z}$ 

```

---

The correctness of the algorithm follows the lemma below. The formal proof is presented in Section 4.3.

**Lemma 4.** *For any tuple  $(s, t, \theta)$  such that  $(s, t)$  is  $\theta$ -extendable,  $\text{EXTPATH}(s, t, \theta)$  outputs a minimum-level  $\theta$ -extendable path from  $s$  to  $t$ .*

By the definition, any  $\theta$ -extendable path from  $f$  to  $t$  becomes a shortest  $\theta$ -alternating path from  $f$  to  $t$  (recall  $\text{level}(\text{out}(f)) = (\infty, \infty)$ ). Hence, it suffices to call  $\text{EXTPATH}(f, g, \text{odd})$  for finding a shortest augmenting path from  $f$  to  $g$ .

### 3.4 Distributed Implementation

In this section, we explain how  $\text{EXTPATH}$  is implemented in the CONGEST model with  $O(\ell)$  rounds, where  $\ell$  represents the length of the path returned by the procedure. We define the distributed implementation of  $\text{EXTPATH}(s, t, \theta)$  as the CONGEST algorithm such that its invocation is triggered by the vertex  $t$ , and finally each vertex in the output path knows its successor and predecessor. As precomputed information, our distributed implementation assumes that each vertex  $u$  locally owns the following knowledge:

- The parent and children of  $u$  in  $T$ , and the value of  $\gamma(u)$ .
- The IDs of two endpoints of  $\text{out}(u)$  (if  $\text{out}(u)$  is not virtual).
- Let  $R_{v, z_v}$  be the unique path from  $v$  to  $z_v$  in  $T$ . If  $u \in V(R_{v, z_v})$ ,  $u$  has the information for routing messages from  $v$  to  $z_v$ . More precisely,  $u$  knows the pair  $(z_v, u')$ , where  $u'$  is the successor of  $u$  in  $R_{v, z_v}$ .

The knowledge above is obtained by the almost same technique as that presented in [19], whose detail is explained in Appendix B. Utilizing the precomputed information stated above, it is very straightforward to implement  $\text{EXTPATH}$  in the CONGEST model. If  $\theta = \gamma(t)$ ,  $t$  sends the triggering message with the information of  $s$  and  $\theta$  to  $\mathbf{p}(t)$  for invoking  $\text{EXTPATH}(s, \mathbf{p}(t), \bar{\theta})$ , as well as adding  $\{\mathbf{p}(t), t\}$  to the output edge set. Otherwise,  $t$  sends the triggering message with  $s$  to the endpoints of  $\text{out}(t)$  along the path  $R_{t, z_t} \circ \text{out}(t)$ . Receiving the message,  $y_t$  and  $z_t$  add  $\text{out}(t)$  to the output edge set, and respectively invoke  $\text{EXTPATH}(s, y_t, \rho(t))$  and  $\text{EXTPATH}(t, z_t, \rho(t))$ . Notice that  $\rho(t)$  is local information available at  $y_t$  and  $z_t$ .

The correctness of this distributed implementation obviously follows Lemma 4. The remaining issue is the running-time analysis. The running time of each invocation is dominated by the rounds for transferring triggering messages. Particularly, the case of  $\theta = \overline{\gamma(t)}$  is the only nontrivial point. One can show that the length of  $|R_{t,z_t}|$  is bounded by the minimum of the lengths of extendable paths from  $s$  to  $y_t$  and from  $t$  to  $z_t$ . The following lemma holds.

**Lemma 5.** *For any  $\overline{\gamma(t)}$ -extendable pair  $(s, t)$ , let  $Z$  be a  $\rho(t)$ -extendable path from  $t$  to  $z_t$  and  $Y$  be a  $\rho(t)$ -extendable path from  $s$  to  $y_t$ . Then,  $|R_{t,z_t}| \leq |Z|$  and  $|R_{t,z_t}| \leq |Y|$  hold.*

Thus, the total running time is bounded as follows:

**Lemma 6.** *Let  $g(h)$  be the worst-case round complexity of  $\text{EXTPATH}(s, t, \theta)$  over all  $(s, t, \theta)$  admitting a  $\theta$ -extendable path from  $s$  to  $t$  of length  $h$ . For any  $0 \leq h \leq \ell$ ,  $g(h) \leq 2h$  holds.*

As we stated in Section 3.3,  $\text{EXTPATH}(f, g, \text{odd})$  suffices to find a shortest augmenting path from  $f$  to  $g$ . For its length  $\ell$ , the running time is at most  $2\ell$ . Putting all the results presented in this section together, we obtain the following corollary.

**Corollary 1.** *There exists an algorithm  $A'(M, \ell)$  satisfying the specification of Lemma 2.*

Combining with Lemmas 1 and 2, Theorem 1 is deduced.

## 4 Detailed Proofs

### 4.1 Fundamental Properties of Alternating Base Trees

Before proceeding to the proofs of key technical lemmas, we present several fundamental properties related to alternating base trees and edge levels. We say that a path  $P$  from  $s$  to  $t$  lies at the inside of  $T_v$  if  $V(P) \subseteq V(T_v)$  holds. We also say that  $P$  lies at the outside of  $T_v$  if  $V(P) \subseteq (V(G) \setminus V(T_v)) \cup \{v\}$ , i.e.,  $P$  is either completely disjoint from  $V(T_v)$  or disjoint from  $V(T_v)$  except for one of its endpoints that is  $v$ .

**Lemma 7.** *The following statements hold:*

- (S1) *For any alternating path  $P$  from  $f$  to  $t$  and an ancestor  $s$  of  $t$  (including the case of  $s = t$ ),  $P$  contains  $\{\mathbf{p}(s), s\}$  if  $\text{level}(P) < \text{level}(\text{out}(s))$ . Hence any alternating path  $P$  from  $f$  to  $t$  not containing  $\{\mathbf{p}(t), t\}$  satisfies  $\text{level}(P) \geq \text{level}(\text{out}(t))$ .*
- (S2) *For any  $t \in V(G)$  and a shortest  $\overline{\gamma(t)}$ -alternating path  $P$  from  $f$  to  $\mathbf{p}(t)$ ,  $P$  does not contain  $t$ .*
- (S3) *The parity of any alternating path  $P$  from  $f$  to  $t$  terminating with edge  $\{\mathbf{p}(t), t\}$  is  $\gamma(t)$ , and thus any  $\overline{\gamma(t)}$ -alternating path from  $f$  to  $t$  does not contain  $\{\mathbf{p}(t), t\}$ .*
- (S4) *For  $u \in V(G)$ ,  $\text{dist}^{\gamma(\mathbf{p}(v))}(\mathbf{p}(v)) < \text{dist}^{\gamma(v)}(v)$  holds, and thus for any descendant  $v$  of  $u$ ,  $\text{dist}^{\gamma(u)}(u) < \text{dist}^{\gamma(v)}(v)$  holds.*
- (S5) *Any shortest  $\gamma(t)$ -alternating  $P$  path from  $f$  to  $t$  lies at the outside of  $T_t$ .*

*Proof. Proof of (S1):* The path  $P$  must contain at least one edge in  $E^*(T_s) \cup \{\{\mathbf{p}(s), s\}\}$  to reach  $s$ . Since only  $\{\mathbf{p}(s), s\}$  has a level less than  $\text{level}(\text{out}(s))$ ,  $P$  necessarily contains it.

**Proof of (S2):** Suppose for contradiction that there exists a shortest  $\overline{\gamma(t)}$ -alternating path  $P$  from  $f$  to  $\mathbf{p}(t)$  that contains  $t$ . Then  $P[f, t]$  is an alternating path from  $f$  to  $t$  of length at

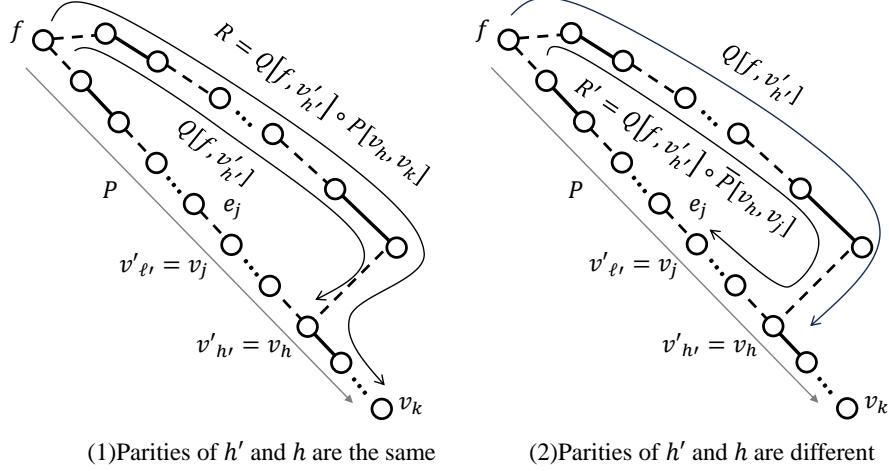


Figure 4: Illustration of Case 2 in the proof of Lemma 8.

most  $|P| - 1 < \text{dist}^{\overline{\gamma(t)}}(f, \mathbf{p}(t))$ . However, by the definition of alternating base trees, the shortest alternating path from  $f$  to  $t$  is of length  $\text{dist}_G^{\overline{\gamma(t)}}(f, \mathbf{p}(t)) + 1$ . It is a contradiction.

**Proof of (S3):** Since any two alternating paths from  $f$  terminating with the same edge must have the same parity, it suffices to show that there exists a  $\gamma(t)$ -alternating path from  $f$  terminating with  $\{\mathbf{p}(t), t\}$ . Let  $Q$  be any shortest  $\overline{\gamma(t)}$ -alternating path from  $f$  to  $\mathbf{p}(t)$ . By statement (S2),  $Q$  does not contain  $t$ , and thus we obtain a  $\gamma(t)$ -alternating path  $P' = Q \circ \{\mathbf{p}(t), t\}$  from  $f$  to  $t$ .

**Proof of (S4):** By the definition of alternating base trees,  $\text{dist}^{\gamma(v)}(v) = \text{dist}^{\overline{\gamma(v)}}(\mathbf{p}(v)) + 1$  holds. Since the case of  $\overline{\gamma(v)} = \gamma(\mathbf{p}(v))$  is obvious, we consider the case of  $\gamma(v) = \gamma(\mathbf{p}(v))$ . By definition, we obtain  $\text{dist}^{\gamma(\mathbf{p}(v))}(\mathbf{p}(v)) < \text{dist}^{\overline{\gamma(\mathbf{p}(v))}}(\mathbf{p}(v))$ . It follows  $\text{dist}^{\gamma(v)}(v) = \text{dist}^{\overline{\gamma(v)}}(\mathbf{p}(v)) + 1 = \text{dist}^{\overline{\gamma(\mathbf{p}(v))}}(\mathbf{p}(v)) + 1 > \text{dist}^{\gamma(\mathbf{p}(v))}(\mathbf{p}(v))$ .

**Proof of (S5):** Every vertex  $u \in V(P)$  satisfies  $\text{dist}^{\gamma(u)}(u) \leq \text{dist}^{\gamma(t)}(t)$ . By statement (S4),  $P$  does not contain any vertex in  $V(T_t) \setminus \{t\}$ , i.e.,  $P$  lies at the outside of  $T_s$ .  $\square$

## 4.2 Proof of Lemma 3

**Lemma 8.** *For any  $\theta \in \{\text{odd}, \text{even}\}$  and  $\theta$ -reachable  $u \in V(G)$ , there exists a shortest  $\theta$ -alternating path  $P$  from  $f$  to  $u$  such that  $\text{level}_{\text{sum}}(P) \leq 2\text{dist}^\theta(u) - 1$ .*

*Proof.* Let  $\mathcal{Q}_k$  be the set of pairs  $(v, \theta)$  such that  $\text{dist}^\theta(v) = k$ . We prove that the lemma holds for any  $k$  and  $(u, \theta) \in \mathcal{Q}_k$  by the induction on  $k$ . (Basis) For any  $(u, \theta) \in \mathcal{Q}_1$ , the edge  $\{f, u\}$  is in the alternating base tree. Since the sum-level of any tree edge is zero, the lemma obviously holds. (Inductive Step) Suppose as the induction hypothesis that the lemma holds for any  $(u', \theta') \in \bigcup_{k' < k} \mathcal{Q}_{k'}$ , and consider the pair  $(u, \theta) \in \mathcal{Q}_k$ . We divide the proof into the following two cases:

**(Case 1)**  $\theta = \gamma(u)$ : Let  $P$  be any minimum-level shortest  $\overline{\gamma(u)}$ -alternating path from  $f$  to  $\mathbf{p}(u)$ . By Lemma 7(S2),  $P$  does not contain  $u$ , and thus we obtain an alternating path  $P' = P \circ \{\mathbf{p}(u), u\}$ , whose parity is  $\gamma(u)$  by Lemma 7(S3). By the definition of alternating base trees, the length of  $P'$  is  $|P| + 1 = \text{dist}^{\overline{\gamma(u)}}(\mathbf{p}(u)) + 1 = \text{dist}^{\gamma(u)}(u) = k$ , i.e.,  $(\mathbf{p}(u), \overline{\gamma(u)}) \in \mathcal{Q}_{k-1}$ . By the induction hypothesis, we obtain  $\text{level}_{\text{sum}}(P) \leq 2k - 3$ . It follows  $\text{level}_{\text{sum}}(P') = \max\{\text{level}_{\text{sum}}(P), \text{level}_{\text{sum}}(\{\mathbf{p}(u), u\})\} \leq 2k - 3$  (recall  $\text{level}_{\text{sum}}(\{\mathbf{p}(u), u\}) = 0$ ). That is, the sum-level of  $P'$  is at most  $2\text{dist}^{\gamma(u)}(u) - 1$ . Since  $|P'| = \text{dist}^{\gamma(u)}(u)$  holds, it is the path satisfying the condition of the lemma.

(**Case 2**)  $\theta = \overline{\gamma(u)}$ : Let  $P = v_0, e_1, v_1, \dots, e_k, v_k$  be any shortest  $\theta$ -alternating path from  $f$  to  $u$ , where  $v_0 = f$  and  $v_k = u$ , and  $e_j$  ( $1 \leq j \leq k$ ) be the edge such that  $\text{level}_{\text{sum}}(e_j) > 2k - 1$  and  $\text{level}_{\text{sum}}(P[v_j, v_k]) \leq 2k - 1$  hold. If such an edge does not exist,  $P$  obviously satisfies  $\text{level}_{\text{sum}}(P) \leq 2\text{dist}^\theta(u) - 1$ , and the proof finishes. Since  $P[f, v_{j-1}]$  is a  $\rho(e_j)$ -alternating path from  $f$  to  $v_{j-1}$ , we obtain  $\text{dist}^{\rho(e_j)}(v_{j-1}) \leq |P[f, v_{j-1}]| = j - 1$ . Then  $\text{dist}^{\rho(e_j)}(v_j) = \text{level}_{\text{sum}}(e_j) - \text{dist}^{\rho(e_j)}(v_{j-1}) > 2k - j \geq k$  holds. Since  $P[f, v_j]$  is a  $\rho(e_j)$ -alternating path from  $f$  to  $v_j$  of length  $j \leq k$ , we have  $\text{dist}^{\rho(e_j)}(v_j) > k \geq \text{dist}^{\overline{\rho(e_j)}}(v_j)$ . That is,  $\overline{\rho(e_j)} = \gamma(v_j)$  holds. It follows  $j < k$  since we assume that the parity  $\theta$  of  $P = P[f, v_k]$  is  $\overline{\gamma(v_k)}$ .

The induction hypothesis provides a shortest  $\gamma(v_j)$ -alternating path  $Q = v'_0, e'_1, v'_1, \dots, e'_\ell, v'_\ell$  (where  $\ell = \text{dist}^{\gamma(v_j)}(v_j)$ ,  $v'_0 = f$ , and  $v'_\ell = v_j$ ) from  $f$  to  $v_j$  of sum-level at most  $2\ell - 1$ . Since  $Q$  is a shortest alternating path from  $f$  to  $v_j$ , its length  $\ell$  is at most  $|P[f, v_j]| = j$ . Let  $h'$  be the smallest index such that  $v'_{h'} \in V(P[v_j, v_k])$ , and  $h$  be the index such that  $v'_h = v_h$  holds ( $1 \leq h' \leq \ell$ ,  $j \leq h \leq k$ ). Note that such  $h'$  and  $h$  necessarily exist because  $v_j$  is contained in both  $V(Q)$  and  $V(P[v_j, v_k])$ .

Consider the case that the parities of  $h$  and  $h'$  are the same (see Fig. 4(1)). Since  $Q[f, v'_{h'}]$  and  $P[v_h, v_k]$  are vertex disjoint except for  $v'_{h'} = v_h$ , we obtain a  $\theta$ -alternating path  $R = Q[f, v'_{h'}] \circ P[v_h, v_k]$  from  $f$  to  $v_k = u$ . The length of  $R$  is  $h' + k - h \leq \ell + k - j \leq k = \text{dist}^\theta(u)$ , i.e.,  $R$  is a shortest  $\theta$ -alternating path. In addition,  $\text{level}_{\text{sum}}(Q) \leq 2j - 1 \leq 2k - 1$  holds and  $P[v_j, v_k]$  does not contain any edge of sum-level greater than  $2k - 1$ . It implies  $\text{level}_{\text{sum}}(R) = \max\{\text{level}_{\text{sum}}(Q[f, v'_{h'}]), \text{level}_{\text{sum}}(P[v_h, v_k])\} \leq 2k - 1$ . That is,  $R$  satisfies the condition of the lemma.

Finally, we rule out the case that the parities of  $h$  and  $h'$  are different (see Fig. 4(2)). Suppose for contradiction that such a case occurs. Then  $R' = Q[f, v'_{h'}] \circ \overline{P}[v_h, v_j]$  becomes a  $\rho(e_j)$ -alternating path from  $f$  to  $v_j$  of length  $h' + h - j \leq j + k - j = k$  (because of  $h' \leq \ell \leq j$  and  $|\overline{P}[v_h, v_j]| \leq |P[v_j, v_k]|$ ). Since  $\text{dist}^{\rho(e_j)}(v_{j-1}) \leq |P[f, v_{j-1}]| \leq k - 1$  holds, we have  $\text{level}_{\text{sum}}(e_j) \leq 2k - 1$  holds. It is a contradiction.  $\square$

**Lemma 9.** For any  $t \in V(G)$  such that  $E^*(T_t)$  is not empty, the following two inequalities hold:

$$(S1) \text{ level}(\text{out}(t)) \leq (\text{dist}^{\text{odd}}(t) + \text{dist}^{\text{even}}(t) - 1, \text{dist}^{\overline{\gamma(t)}}(t) - 1).$$

$$(S2) \text{ dist}^{\overline{\gamma(t)}}(t) \geq \text{dist}^{\rho(t)}(z_t) + \text{dist}^{\rho(t)}(y_t) + 1 - \text{dist}^{\gamma(t)}(t).$$

*Proof.* Let  $P$  be any shortest  $\gamma(t)$ -alternating path from  $f$  to  $t$ . By Lemma 7(S5),  $P$  lies at the outside of  $T_t$ . Let  $Q$  be any shortest  $\gamma(t)$ -alternating path from  $f$  to  $t$ . By Lemma 7(S3),  $Q$  does not contain  $\{p(t), t\}$  and thus contains at least one edge in  $E^*(T_t)$ . Let  $e = \{y', z'\}$ , where  $z' \in V(T_t)$ , be the edge in  $E(Q) \cap E^*(T_t)$  such that  $Q[z', t]$  does not contain any edge in  $E^*(T_t)$ . Since  $\overline{Q}[t, z']$  lies at the inside of  $T_t$ ,  $P \circ \overline{Q}[t, z']$  is a  $\rho(e)$ -alternating path from  $f$  to  $z'$  and thus  $\text{dist}^{\rho(e)}(z') \leq |P \circ \overline{Q}[t, z']|$  holds. In addition, since  $Q[f, y']$  is a  $\rho(e)$ -alternating path from  $f$  to  $y'$ ,  $\text{dist}^{\rho(e)}(y') \leq |Q[f, y']|$  also holds.

**Proof of (S1):** Since  $\text{level}(\text{out}(t)) \leq \text{level}(e)$  holds, it suffices to show  $\text{level}(e) \leq (\text{dist}^{\text{odd}}(t) + \text{dist}^{\text{even}}(t) - 1, \text{dist}^{\overline{\gamma(t)}}(t) - 1)$ . We have

$$\begin{aligned} \text{level}_{\text{sum}}(e) &= \text{dist}^{\rho(e)}(z') + \text{dist}^{\rho(e)}(y') \\ &\leq |P \circ \overline{Q}[t, z']| + |Q[f, y']| \\ &= |P| + |\overline{Q}[t, z']| + |Q[f, y']| \\ &= |P| + |Q| - 1 \end{aligned}$$

$$\begin{aligned}
&= \text{dist}^{\gamma(t)}(t) + \text{dist}^{\overline{\gamma(t)}}(t) - 1 \\
&= \text{dist}^{\text{odd}}(t) + \text{dist}^{\text{even}}(t) - 1.
\end{aligned}$$

Next, we upper bound  $\text{level}_{\max}(e) = \max\{\text{dist}^{\rho(e)}(y'), \text{dist}^{\rho(e)}(z')\}$  by  $\text{dist}^{\overline{\gamma(t)}}(t) - 1$ . We have  $\text{dist}^{\rho(e)}(y') \leq |Q[f, y']| \leq |Q| - 1 = \text{dist}^{\overline{\gamma(t)}}(t) - 1$ . Since  $z'$  is a descendant of  $t$  (including  $t$  itself),  $|Q[f, z']| > \text{dist}^{\gamma(t)}(t)$  holds as follows. If  $z' = t$ , then  $|Q[f, z']| = |Q| > |P| = \text{dist}^{\gamma(t)}(t)$ ; otherwise,  $|Q[f, z']| \geq \text{dist}^{\gamma(z')}(z') > \text{dist}^{\gamma(t)}(t)$  by Lemma 7(S4). Then  $\text{dist}^{\rho(e)}(z') \leq |P \circ \overline{Q}[t, z']| = \text{dist}^{\gamma(t)}(t) + (\text{dist}^{\overline{\gamma(t)}}(t) - |Q[f, z']|) < \text{dist}^{\overline{\gamma(t)}}(t)$  holds. Consequently, we have  $\text{level}_{\max}(e) \leq \text{dist}^{\overline{\gamma(t)}}(t) - 1$ .

**Proof of (S2):** Since  $P \circ \overline{Q}[t, z']$  is a  $\rho(e)$ -alternating path from  $f$  to  $z'$ ,  $\text{dist}^{\rho(e)}(z') \leq |P \circ \overline{Q}[t, z']| = |Q[z', t]| + \text{dist}^{\gamma(t)}(t)$  holds, i.e.,  $|Q[z', t]| \geq \text{dist}^{\rho(e)}(z') - \text{dist}^{\gamma(t)}(t)$  holds. Then we have

$$\begin{aligned}
\text{dist}^{\overline{\gamma(t)}}(t) &= |Q| \\
&= |Q[f, y'] \circ e \circ Q[z', t]| \\
&\geq \text{dist}^{\rho(e)}(y') + 1 + \text{dist}^{\rho(e)}(z') - \text{dist}^{\gamma(t)}(t) \\
&\geq \text{dist}^{\rho(t)}(y_t) + 1 + \text{dist}^{\rho(t)}(z_t) - \text{dist}^{\gamma(t)}(t),
\end{aligned}$$

where the last inequality follows from the definition of MOEs, i.e.,  $\text{level}_{\text{sum}}(e) \geq \text{level}_{\text{sum}}(\text{out}(t))$  implies that  $\text{dist}^{\rho(e)}(z') + \text{dist}^{\rho(e)}(y') \geq \text{dist}^{\rho(t)}(z_t) + \text{dist}^{\rho(t)}(y_t)$ .  $\square$

**Lemma 10.** *For any  $t \in V(G)$  such that  $E^*(T_t)$  is not empty, if  $\text{dist}^{\rho(t)}(y_t) \leq \text{dist}^{\rho(t)}(z_t)$  holds,  $\text{dist}^{\overline{\rho(t)}}(z_t) \leq \text{dist}^{\rho(t)}(y_t) + 1$  holds.*

*Proof.* Utilizing Lemma 8 with  $\theta = \rho(t)$  and  $u = y_t$ , there exists a shortest  $\rho(t)$ -alternating path  $P$  from  $f$  to  $y_t$  such that  $\text{level}_{\text{sum}}(P) \leq 2\text{dist}^{\rho(t)}(y_t) - 1$  (note that  $y_t$  is  $\rho(t)$ -reachable because any edge  $e \in E^*(T_t)$  has a finite edge level and thus  $\text{dist}^{\rho(t)}(y_t)$  is finite). Since we assume  $\text{dist}^{\rho(t)}(y_t) \leq \text{dist}^{\rho(t)}(z_t)$  we also have  $\text{level}_{\text{sum}}(\text{out}(t)) = \text{dist}^{\rho(t)}(y_t) + \text{dist}^{\rho(t)}(z_t) \geq 2\text{dist}^{\rho(t)}(y_t)$ , i.e., no outgoing edge of  $T_t$  is contained in  $P$ . It implies that  $P$  lies at the outside of  $T_t$  and thus does not contain  $z_t$ . Hence  $P' = P \circ \{y_t, z_t\}$  is a  $\rho(t)$ -alternating path from  $f$  to  $z_t$ , whose length is  $\text{dist}^{\rho(t)}(y_t) + 1$ . That is,  $\text{dist}^{\overline{\rho(t)}}(z_t) \leq \text{dist}^{\rho(t)}(y_t) + 1$  holds.  $\square$

**Lemma 11.** *For any  $t, t' \in V(G)$  such that  $E^*(T_t)$  and  $E^*(T_{t'})$  are not empty, if  $\text{level}(\text{out}(t)) = \text{level}(\text{out}(t'))$  holds, then  $\rho(t) = \rho(t')$  holds.*

*Proof.* In particular, we have  $\text{level}_{\max}(\text{out}(t)) = \text{level}_{\max}(\text{out}(t'))$ , i.e.,  $\max\{\text{dist}^{\rho(t)}(y_t), \text{dist}^{\rho(t)}(z_t)\} = \max\{\text{dist}^{\rho(t')}(y_{t'}), \text{dist}^{\rho(t')}(z_{t'})\}$ . The both sides must have the same parity, which means  $\rho(t) = \rho(t')$ .  $\square$

**Lemma 3 (Repeat).** *For any  $t \in V(G)$  such that  $E^*(T_t) \neq \emptyset$ , the following three statements hold (see Fig. 3):*

- (S1) *There exists a shortest  $\rho(t)$ -alternating path  $Y'$  from  $f$  to  $y_t$  of level less than  $\text{level}(\text{out}(t))$ .*
- (S2) *There exists a shortest  $\rho(t)$ -alternating path  $Z'$  from  $f$  to  $z_t$  of level less than  $\text{level}(\text{out}(t))$ .*
- (S3) *The path  $X' = Y' \circ \{y_t, z_t\} \circ \overline{Z}'[z_t, t]$  is a shortest  $\overline{\gamma(t)}$ -alternating path from  $f$  to  $t$ . This does not depend on the choices of  $Y'$  and  $Z'$  in the above two statements.*

*Proof.* For  $k \in \mathbb{N}^2$ , let  $\mathcal{V}_k$  be the set of the vertices  $v$  such that  $E^*(T_v)$  is not empty and  $\text{level}(\text{out}(v)) = k$  holds. We show that the lemma holds for any  $k$  and  $t \in \mathcal{V}_k$  following the induction on  $k$ . The base case of  $k = (0, 0)$  is obvious (as  $\mathcal{V}_k = \emptyset$ ), and thus we focus on the inductive step. The proof of the inductive step consists of the following three parts:



(P1) Supposing as the induction hypothesis that the lemma holds for any  $t' \in \bigcup_{k' < k} \mathcal{V}_{k'}$ , show the statement (S2) holds for any vertex  $t \in \mathcal{V}_k$ .

(P2) Supposing that the lemma holds for any  $t' \in \bigcup_{k' < k} \mathcal{V}_{k'}$  as the induction hypothesis, and that the statement (S2) holds for any  $t \in \mathcal{V}_k$  by the consequence of (P1), show the statement (S1) holds for any  $t \in \mathcal{V}_k$ .

(P3) For any  $t \in \mathcal{V}_k$ , (S1) and (S2) imply (S3).

**The proof of (P1):** The proof is divided into the following three cases:

**(Case 1)**  $\text{dist}^{\rho(t)}(z_t) \leq \text{dist}^{\rho(t)}(y_t)$ : Then  $\text{level}_{\text{sum}}(\text{out}(t)) \geq 2\text{dist}^{\rho(t)}(z_t)$  holds. By Lemma 8, there exists a shortest  $\rho(t)$ -alternating path  $P$  from  $f$  to  $z_t$  of sum-level at most  $2\text{dist}^{\rho(t)}(z_t) - 1 < \text{level}_{\text{sum}}(\text{out}(t))$ , i.e.,  $P$  is the path satisfying the condition of (S2).

**(Case 2)**  $\text{dist}^{\rho(t)}(z_t) > \text{dist}^{\rho(t)}(y_t)$  and  $\text{level}(\text{out}(z_t)) < \text{level}(\text{out}(t))$ : By the induction hypothesis, there exists a shortest  $\overline{\gamma(z_t)}$ -alternating path  $X'$  from  $f$  to  $z_t$  whose level is bounded by  $\text{level}(\text{out}(z_t)) < \text{level}(\text{out}(t))$ . Since  $\text{dist}^{\rho(t)}(z_t) > \text{dist}^{\rho(t)}(y_t)$ , by Lemma 10, we obtain  $\text{dist}^{\overline{\rho(t)}}(z_t) \leq \text{dist}^{\rho(t)}(y_t) + 1 \leq \text{dist}^{\rho(t)}(z_t)$ . It implies  $\overline{\rho(t)} = \gamma(z_t)$  because of  $\text{dist}^{\overline{\rho(t)}}(z_t) \neq \text{dist}^{\rho(t)}(z_t)$ . Then  $X'$  is the path satisfying the condition of (S1).

**(Case 3)**  $\text{dist}^{\rho(t)}(z_t) > \text{dist}^{\rho(t)}(y_t)$  and  $\text{level}(\text{out}(z_t)) \geq \text{level}(\text{out}(t))$ : We show that this case never happens by leading to a contradiction. Let  $t' = z_t$ ,  $y_{t'} = y_{z_t}$ , and  $z_{t'} = z_{z_t}$  for short. Since  $\text{out}(t)$  is an edge incident to  $z_t$  which reaches the vertex  $y_t$  at the outside of  $T_t$ ,  $\text{out}(t)$  is also an outgoing edge of  $T_{z_t}$ , i.e.,  $\text{level}(\text{out}(z_t)) \leq \text{level}(\text{out}(t))$  holds. Combining with the assumption of  $\text{level}(\text{out}(z_t)) \geq \text{level}(\text{out}(t))$ , we obtain  $\text{level}(\text{out}(z_t)) = \text{level}(\text{out}(t))$ . In addition, by the rule of choosing canonical MOEs, the edges of level  $\text{level}(\text{out}(z_t))$  incident to  $z_t$  have higher priority. Hence  $\text{out}(z_t)$  is incident to  $z_t$ , i.e.,  $z_{t'} = t'$  holds. Since  $\rho(t) = \rho(t')$  is obtained by Lemma 11, the assumption of  $\text{dist}^{\rho(t)}(z_t) > \text{dist}^{\rho(t)}(y_t)$  implies  $\text{dist}^{\rho(t')}(z_{t'}) > \text{dist}^{\rho(t')}(y_{t'})$ . Then we obtain  $\text{dist}^{\overline{\rho(t')}}(z_{t'}) \leq \text{dist}^{\rho(t')}(y_{t'}) + 1$  by Lemma 10. Putting all the (in)equalities above together, we conclude  $\text{dist}^{\overline{\rho(t')}}(t') \leq \text{dist}^{\rho(t')}(y_{t'}) + 1 \leq \text{dist}^{\rho(t')}(t')$ . Then, the level of  $\text{out}(t')$  is bounded as follows:

$$\begin{aligned} \text{level}_{\text{sum}}(\text{out}(t')) &= \text{dist}^{\rho(t')}(y_{t'}) + \text{dist}^{\rho(t')}(z_{t'}) \\ &\geq \text{dist}^{\overline{\rho(t')}}(t') - 1 + \text{dist}^{\rho(t')}(t') \\ &= \text{dist}^{\text{odd}}(t') + \text{dist}^{\text{even}}(t') - 1, \\ \text{level}_{\text{max}}(\text{out}(t')) &= \max\{\text{dist}^{\rho(t')}(y_{t'}), \text{dist}^{\rho(t')}(z_{t'})\} \\ &= \text{dist}^{\rho(t')}(t') \\ &> \text{dist}^{\overline{\gamma(t')}}(t') - 1. \end{aligned}$$

That is,  $\text{level}(\text{out}(t')) > (\text{dist}^{\text{odd}}(t') + \text{dist}^{\text{even}}(t') - 1, \text{dist}^{\overline{\gamma(t')}}(t') - 1)$ . It contradicts Lemma 9(S1).

**The proof of (P2):** The proof is divided into the two cases below:

**(Case 1)**  $\text{level}(\text{out}(t)) \leq \text{level}(\text{out}(y_t))$ : Suppose that  $y_t$  is an ancestor of  $z_t$ . If  $\rho(t) = \gamma(y_t)$ , then a shortest  $\rho(t)$ -alternating path from  $f$  to  $y_t$  plus  $\text{out}(t)$  is a  $\rho(t)$ -alternating path from  $f$  to  $z_t$ , say  $P$ . By Lemma 7(S4), we obtain  $\text{dist}^{\gamma(z_t)}(z_t) \leq \text{dist}^{\rho(t)}(z_t) \leq |P| = \text{dist}^{\gamma(y_t)}(y_t) + 1 \leq \text{dist}^{\gamma(z_t)}(z_t)$ , which implies  $\rho(t) = \gamma(z_t)$ . However, this implies that  $\text{dist}^{\gamma(y_t)}(y_t)$  and  $\text{dist}^{\gamma(z_t)}(z_t)$  have the same parity, so  $\text{dist}^{\gamma(y_t)}(y_t) \leq \text{dist}^{\gamma(z_t)}(z_t) - 2$  holds. Thus we conclude  $\text{dist}^{\gamma(z_t)}(z_t) \leq |P| = \text{dist}^{\gamma(y_t)}(y_t) + 1 \leq \text{dist}^{\gamma(z_t)}(z_t) - 1$ , a contradiction.

Otherwise,  $\rho(t) = \overline{\gamma(y_t)}$ . By Lemma 7(S4) again, we have

$$\begin{aligned} \text{level}_{\text{sum}}(\text{out}(t)) &= \text{dist}^{\rho(t)}(y_t) + \text{dist}^{\rho(t)}(z_t) \\ &\geq \text{dist}^{\overline{\gamma(y_t)}}(y_t) + \text{dist}^{\gamma(z_t)}(z_t) \\ &> \text{dist}^{\overline{\gamma(y_t)}}(y_t) + \text{dist}^{\gamma(y_t)}(y_t), \end{aligned}$$

which implies that there exists an edge  $e \in T^*(y_t)$  such that  $\text{level}_{\text{sum}}(\text{out}(t)) > \text{level}_{\text{sum}}(e)$  (since the union of shortest  $\gamma(y_t)$ - and  $\overline{\gamma(y_t)}$ -alternating paths from  $f$  to  $y_t$  must contain an edge in  $T^*(y_t)$  of sum-level at most  $\text{dist}^{\overline{\gamma(y_t)}}(y_t) + \text{dist}^{\gamma(y_t)}(y_t)$ ). Hence  $\text{level}(\text{out}(t)) > \text{level}(\text{out}(y_t))$ , a contradiction.

Now we assume that  $y_t$  is not an ancestor of  $z_t$ , which means that  $\text{out}(t)$  is also an outgoing edge of  $T_{y_t}$ . Hence  $\text{level}(\text{out}(t)) \geq \text{level}(\text{out}(y_t))$  necessarily holds and it suffices to consider the case of  $\text{level}(\text{out}(t)) = \text{level}(\text{out}(y_t))$ , i.e.,  $y_t \in \mathcal{V}_k$ . Since  $\text{out}(t)$  is the edge incident to  $y_t$ ,  $\text{out}(y_t)$  must be incident to  $y_t$  by the rule of choosing canonical MOEs. By applying (S1) of this lemma (recall that the proof of (P2) assumes that (P1) holds even for  $t \in \mathcal{V}_k$ ), we have a shortest  $\rho(y_t)$ -alternating path  $Z'$  from  $f$  to  $y_t$  of level less than  $\text{level}(\text{out}(y_t)) = \text{level}(\text{out}(t))$ . By Lemma 11,  $\text{level}(\text{out}(t)) = \text{level}(\text{out}(y_t))$  implies  $\rho(t) = \rho(y_t)$ . Hence  $Z'$  is the one satisfying the condition of (S1).

**(Case 2)  $\text{level}(\text{out}(t)) > \text{level}(\text{out}(y_t))$ :** In this case,  $y_t \in \bigcup_{k' < k} \mathcal{V}_{k'}$  holds. By the induction hypothesis, we obtain a shortest  $\rho(y_t)$ -alternating path  $Z'$  from  $f$  to  $z_{y_t}$  of level less than  $\text{level}(\text{out}(y_t)) < \text{level}(\text{out}(t))$ . Since  $\text{level}(Z') < \text{level}(\text{out}(y_t))$  holds,  $Z'$  must contain  $\{\mathfrak{p}(y_t), y_t\}$  to reach the inside of  $T_{y_t}$ , and contains no outgoing edge in  $E^*(T_{y_t})$ . Then  $Z'[f, y_t]$  lies at the outside of  $T_{y_t}$  and  $Z'[y_t, z_{y_t}]$  lies at the inside of  $T_{y_t}$ . It implies that  $Z'[f, y_t]$  is a shortest  $\gamma(y_t)$ -alternating path from  $f$  to  $y_t$  because if there exists a  $\gamma(y_t)$ -alternating path  $Q$  from  $f$  to  $y_t$  shorter than  $Z'[f, y_t]$ ,  $Q \circ Z'[y_t, z_{y_t}]$  becomes a  $\rho(y_t)$ -alternating path shorter than  $Z'$  (note that  $Q$  lies at the outside of  $T_{y_t}$  by Lemma 7(S5) and thus  $Q \circ Z'[y_t, z_{y_t}]$  necessarily becomes an alternating path). In addition, by the induction hypothesis, we also has a shortest  $\overline{\gamma(y_t)}$ -alternating path  $X'$  from  $f$  to  $y_t$  of level less than  $\text{level}(\text{out}(y_t)) < \text{level}(\text{out}(t))$ . Consequently, either  $Z'[f, y_t]$  or  $X'$  has parity  $\rho(t)$ , which is the path satisfying the condition of (S1).

**The proof of (P3):** Suppose the existence of  $Y'$  and  $Z'$  stated in (S1) and (S2) for  $t \in \mathcal{V}_k$ . Since the levels of  $Y'$  and  $Z'$  are less than  $\text{out}(t)$ , they respectively lie at the outside of  $T_t$  and at the inside of  $T_t$ . Hence the path  $X' = Y' \circ \text{out}(t) \circ \overline{Z'}[z_t, t]$  is an alternating path from  $f$  to  $t$ . Since  $Z'[f, t]$  is an alternating path of level less than  $\text{level}(\text{out}(t))$ , by Lemma 7(S1) and (S4),  $Z'[f, t]$  is a shortest  $\gamma(t)$ -alternating path from  $f$  to  $t$ . Hence  $X'$  is a  $\overline{\gamma(t)}$ -alternating path, and we have  $|\overline{Z'}[z_t, t]| \leq |Z'| - |Z'[f, t]| = \text{dist}^{\rho(t)}(z_t) - \text{dist}^{\gamma(t)}(t)$ . Also,  $|Y'| = \text{dist}^{\rho(t)}(y_t)$  holds. By Lemma 9(S2),  $\text{dist}^{\overline{\gamma(t)}}(t) \geq \text{dist}^{\rho(t)}(z_t) + \text{dist}^{\rho(t)}(y_t) + 1 - \text{dist}^{\gamma(t)}(t)$  holds. Combining the three (in)equalities, we conclude  $|X'| = |Y'| + 1 + |\overline{Z'}[z_t, t]| \leq \text{dist}^{\overline{\gamma(t)}}(t)$ .  $\square$

### 4.3 Correctness of EXTPATH

For the proof of Lemma 4, we present two auxiliary lemmas.

**Lemma 12.** *For any  $\theta$ -extendable pair  $(s, t)$ , every  $\theta$ -extendable path from  $s$  to  $t$  has the same length.*

*Proof.* Let  $P$  and  $Q$  be two different  $\theta$ -extendable paths from  $s$  to  $t$ , and  $P'$  and  $Q'$  be their  $\theta$ -extensions. Since the levels of  $P'$  and  $Q'$  are both less than  $\text{out}(s)$ , they contain  $\{\mathfrak{p}(s), s\}$  and do not contain any outgoing edge of  $T_s$ . It implies that  $P'[f, \mathfrak{p}(s)]$  and  $Q'[f, \mathfrak{p}(s)]$  do not contain any vertex in  $T_s$ , and  $P'[s, t]$  and  $Q'[s, t]$  consist only of the vertices in  $T_s$ . Suppose for contradiction

that  $P = P'[s, t]$  and  $Q = Q'[s, t]$  have different lengths. Without loss of generality, we suppose  $|P'[s, t]| < |Q'[s, t]|$ . Since  $|P'| = |Q'|$  holds, we have  $|P'[f, \mathbf{p}(s)]| > |Q'[f, \mathbf{p}(s)]|$ . It implies that the path  $Q'[f, \mathbf{p}(s)] \circ \{\mathbf{p}(s), s\} \circ Q'[s, t]$  becomes a  $\theta$ -alternating path shorter than  $P'$  and  $Q'$ . It is a contradiction.  $\square$

For each  $(s, t, \theta)$  such that  $(s, t)$  is  $\theta$ -extendable, pick a minimum-level  $\theta$ -extendable path  $X_{s,t,\theta}$  from  $s$  to  $t$  (with any tie-breaking rule). Let  $\mathcal{Q}_i$  be the set of all tuples  $(s, t, \theta)$  such that  $|X_{s,t,\theta}| = i$ . Note that  $\mathcal{Q}_i$  to which  $(s, t, \theta)$  belongs is not affected by the tie-breaking rule for the choice of  $X_{s,t,\theta}$  because of Lemma 12.

**Lemma 13.** *For any  $i$  and  $(s, t, \gamma(t)) \in \mathcal{Q}_i$ ,  $\{\mathbf{p}(t), t\} \in E(X_{s,t,\gamma(t)})$  holds.*

*Proof.* If  $E^*(T_t)$  is empty,  $\{\mathbf{p}(t), t\}$  is a bridge of  $G$  separating  $s$  and  $t$ , and thus  $X_{s,t,\gamma(t)}$  necessarily contains  $\{\mathbf{p}(t), t\}$ . Suppose that  $E^*(T_t) \neq \emptyset$  and for contradiction that  $\{\mathbf{p}(t), t\} \notin E(X_{s,t,\gamma(t)})$ . Then  $X_{s,t,\gamma(t)}$  must contain an edge in  $E^*(T_t)$  to reach vertices in  $T_t$ . Then we have  $\text{level}(X_{s,t,\gamma(t)}) \geq \text{out}(t)$ . Since  $\text{level}(X_{s,t,\gamma(t)}) \leq \text{level}(W) < \text{level}(\text{out}(s))$  by definition, where  $W$  is any  $\gamma(t)$ -extension of  $X_{s,t,\gamma(t)}$ , it follows  $\text{level}(\text{out}(t)) < \text{level}(\text{out}(s))$ . Lemma 3(S2) states that there exists an alternating path  $Z'$  from  $f$  to  $z_t$  of level less than  $\text{level}(\text{out}(t))$ . Since  $\text{level}(Z') < \text{level}(\text{out}(t)) < \text{level}(\text{out}(s))$  holds and  $s$  and  $t$  are the ancestors of  $z_t$ , both  $s$  and  $t$  are contained in  $Z'$  by Lemma 7(S1). By lemma 7(S3), any alternating path from  $f$  to  $t$  terminating with edge  $\{\mathbf{p}(t), t\}$  is a  $\gamma(t)$ -alternating path. It follows that the subpath of  $Z'$  from  $s$  to  $t$  is a  $\gamma(t)$ -extendable path of level less than  $\text{level}(\text{out}(t)) \leq \text{level}(X_{s,t,\gamma(t)})$ , which contradicts the definition that  $X_{s,t,\gamma(t)}$  is the minimum-level.  $\square$

Utilizing the lemmas above, we prove Lemma 4.

**Lemma 4 (Repeat).** *For any tuple  $(s, t, \theta)$  such that  $(s, t)$  is  $\theta$ -extendable,  $\text{EXTPATH}(s, t, \theta)$  outputs a minimum-level  $\theta$ -extendable path from  $s$  to  $t$ .*

*Proof.* We prove the lemma for any  $i$  and  $(s, t, \theta) \in \mathcal{Q}_i$  following the induction on  $i$ . In the following argument, let  $Y$  and  $Z$  be the paths computed in  $\text{EXTPATH}(s, t, \theta)$  (see Algorithm 1), and  $W$  be any  $\theta$ -extension of  $X_{s,t,\theta}$ .

**(Basis)** For  $i = 0$ , the algorithm returns the path of length zero, and thus the lemma obviously holds.

**(Inductive step)** Suppose as the induction hypothesis that the lemma holds for any  $i' < i$  and  $(s', t', \theta') \in \mathcal{Q}_{i'}$ , and consider a tuple  $(s, t, \theta) \in \mathcal{Q}_i$ . We divide the proof into two cases according to the value of  $\theta$ :

**(Case 1)**  $\theta = \gamma(t)$ : First, we prove  $(s, \mathbf{p}(t), \bar{\theta}) \in \mathcal{Q}_{i-1}$ . By Lemma 13,  $\mathbf{p}(t)$  is on the path  $X_{s,t,\theta}$ , and thus on the path  $W$ . Then  $W[f, \mathbf{p}(t)]$  is a shortest  $\bar{\theta}$ -alternating path from  $f$  to  $\mathbf{p}(t)$  of  $\text{level}(W) < \text{level}(\text{out}(s))$  because of the equality  $|W[f, \mathbf{p}(t)]| = |W| - 1 = \text{dist}^\theta(t) - 1 = \text{dist}^{\bar{\theta}}(\mathbf{p}(t))$  derived from the definition of alternating base trees. It implies that  $W[f, \mathbf{p}(t)]$  is an  $\bar{\theta}$ -extension of  $(s, \mathbf{p}(t))$ , i.e.,  $(s, \mathbf{p}(t))$  is  $\bar{\theta}$ -extendable. In addition,  $X_{s,t,\theta}[s, \mathbf{p}(t)]$  is a  $\bar{\theta}$ -extendable path from  $s$  to  $\mathbf{p}(t)$  of length  $|X_{s,t,\theta}| - 1$ , i.e.,  $(s, \mathbf{p}(t), \bar{\theta}) \in \mathcal{Q}_{i-1}$ .

By the induction hypothesis,  $Y = \text{EXTPATH}(s, \mathbf{p}(t), \bar{\theta})$  is a minimum-level  $\bar{\theta}$ -extendable path from  $s$  to  $\mathbf{p}(t)$ . By Lemma 7(S2),  $Y$  does not contain  $t$ . Thus  $P = Y \circ \{\mathbf{p}(t), t\} = \text{EXTPATH}(s, t, \theta)$  is an alternating path from  $s$  to  $t$ . It is necessarily a minimum-level  $\theta$ -extendable path: We have already shown that  $X_{s,t,\theta}[s, \mathbf{p}(t)]$  is a  $\bar{\theta}$ -extendable path from  $s$  to  $\mathbf{p}(t)$ . Since  $Y$  is the minimum-level,  $\text{level}(Y) \leq \text{level}(X_{s,t,\theta}[s, \mathbf{p}(t)])$  holds. The levels of  $X_{s,t,\theta}$  and  $P$  are respectively decided by the levels of their prefixes  $X_{s,t,\theta}[s, \mathbf{p}(t)]$  and  $P[s, \mathbf{p}(t)] = Y$  (recall that any edge in  $T$  has level  $(0, 0)$ ). Hence we can conclude  $\text{level}(X_{s,t,\theta}) = \text{level}(X_{s,t,\theta}[s, \mathbf{p}(t)]) \geq \text{level}(Y) = \text{level}(P)$ . Since

$\text{level}(P) < \text{level}(\text{out}(s))$  holds,  $P$  lies at the inside of  $T_s$ . Then  $P' = W[f, s] \circ P$  becomes an alternating path of length  $|W|$  because  $|Y| = |W[s, p(t)]|$  holds by Lemma 12, i.e.,  $P$  admits a  $\theta$ -extension  $P'$  and thus is a  $\theta$ -extendable path.

**(Case 2)**  $\theta = \overline{\gamma(t)}$ : By Lemma 7(S3),  $W$  does not contain  $\{p(t), t\}$ . Applying Lemma 7(S1), we obtain  $\text{level}(\text{out}(t)) \leq \text{level}(W) < \text{level}(\text{out}(s))$ . Let  $X'$ ,  $Y'$ , and  $Z'$  be the alternating paths defined in Lemma 3. Two alternating paths  $Y' \circ \{y_t, z_t\}$  and  $Z'$  from  $f$  to  $z_t$  have levels at most  $\text{level}(\text{out}(t)) < \text{level}(\text{out}(s))$ . Since  $s$  is an ancestor of  $t$  and thus that of  $z_t$ ,  $s \in V(Y')$  holds by Lemma 7(S1). Then  $Y'$  and  $Z'$  are respectively  $\rho(t)$ -extensions of  $(s, y_t)$  and  $(t, z_t)$ , and  $Y'[s, y_t]$  and  $Z'[t, z_t]$  are the corresponding  $\rho(t)$ -extendable paths. Since  $X'$  is also a  $\theta$ -extension of  $(s, t)$ ,  $X'[s, t]$  is a  $\theta$ -extendable path from  $s$  to  $t$ . It follows  $|Y'[s, y_t]| + |Z'[t, z_t]| = |X_{s, y_t, \rho(t)}| + |X_{t, z_t, \rho(t)}| < |X'[s, t]| = |X_{s, t, \theta}|$ , where the first and last equalities come from Lemma 12. It implies  $|Y'[s, y_t]| < |X_{s, t, \theta}|$  and  $|Z'[t, z_t]| < |X_{s, t, \theta}|$ , i.e.,  $(s, y_t), (t, z_t) \in \bigcup_{i' < i} \mathcal{Q}_{i'}$ . By the induction hypothesis,  $Y = \text{EXTPATH}(s, y_t, \rho(t))$  and  $Z = \text{EXTPATH}(t, z_t, \rho(t))$  are respectively the minimum-level  $\rho(t)$ -alternating paths from  $s$  to  $y_t$  and from  $t$  to  $z_t$ . Since  $Y$  and  $Z$  are the minimum-level, their levels are respectively upper bounded by  $\text{level}(Y'[s, y_t])$  and  $\text{level}(Z'[t, z_t])$ , and here, less than  $\text{level}(\text{out}(t))$ . Since we have already shown  $\text{level}(\text{out}(t)) < \text{level}(\text{out}(s))$ , it also follows  $\text{level}(Y) < \text{level}(\text{out}(s))$  and  $\text{level}(Z) < \text{level}(\text{out}(s))$ . Then  $Y$  and  $Z$  do not contain any edge in  $E^*(T_s) \cup E^*(T_t)$ , i.e.,  $Y$  lies at the inside of  $T_s$  but the outside of  $T_t$ , and  $Z$  lies at the inside of  $T_t$ . Consequently,  $P = Y \circ \{y_t, z_t\} \circ \bar{Z}$  returned by the algorithm is an alternating path of level  $\text{level}(\text{out}(t))$  lying at the inside of  $T_s$ . By Lemma 12, we have

$$\begin{aligned} \text{dist}^\theta(t) &= |X'| \\ &= |X'[f, s]| + |X'[s, y_t]| + 1 + |X'[z_t, t]| \\ &= |X'[f, s]| + |Y| + |Z| + 1 \\ &= |X'[f, s]| + |P|. \end{aligned}$$

That is,  $X'[f, s] \circ P$  is a shortest  $\theta$ -alternating path from  $f$  to  $t$ . The level of  $X'[f, s] \circ P$  is bounded as follows:

$$\begin{aligned} \text{level}(X'[f, s] \circ P) &= \max\{\text{level}(X'[f, s]), \text{level}(P)\} \\ &\leq \max\{\text{level}(X'), \text{level}(Y), \text{level}(Z), \text{level}(\{y_t, z_t\})\} \\ &= \text{level}(\text{out}(t)) \\ &< \text{level}(\text{out}(s)), \end{aligned}$$

where we already see  $\text{level}(X') = \text{level}(\{y_t, z_t\}) = \text{level}(\text{out}(t))$ ,  $\text{level}(Y) \leq \text{level}(\text{out}(t))$ , and  $\text{level}(Z) < \text{level}(\text{out}(t))$ . It implies  $X'[f, s] \circ P$  is a  $\theta$ -extension of  $(s, t)$ , and thus  $P$  is a  $\theta$ -extendable path of level at most  $\text{level}(\text{out}(t))$ . Since  $W$  does not contain  $\{p(t), t\}$  and  $s$  is at the outside of  $T_t$ ,  $W[s, t] = X_{s, t, \theta}$  must contain at least one outgoing edge of  $T_t$ . Then  $\text{level}(X_{s, t, \theta}) \geq \text{level}(\text{out}(t))$  holds and thus the constructed path  $P$  is of the minimum-level.  $\square$

#### 4.4 Running-Time Analysis of Distributed Implementation

**Lemma 5** (Repeat). *For any  $\overline{\gamma(t)}$ -extendable pair  $(s, t)$ , let  $Z$  be a  $\rho(t)$ -extendable path of  $(t, z_t)$  and  $Y$  be a  $\rho(t)$ -extendable path of  $(s, y_t)$ . Then  $|R_{t, z_t}| \leq |Z|$  and  $|R_{t, z_t}| \leq |Y|$  hold.*

*Proof.* By Lemma 7(S4),  $\text{dist}^{\gamma(t)}(t) \geq \text{dist}^{\rho(t)}(p(t)) + 1$  holds for any  $t \in V(G) \setminus \{f\}$ . It implies  $|R_{t, z_t}| \leq \text{dist}^{\gamma(z_t)}(z_t) - \text{dist}^{\gamma(t)}(t)$ . Let  $Z'$  be a shortest  $\rho(t)$ -alternating path from  $f$  to  $z_t$  of level less than  $\text{level}(\text{out}(t))$  by Lemma 3. Then  $|Z'| \geq \text{dist}^{\gamma(z_t)}(z_t)$  holds. Since  $Z'[t, z_t]$  is a  $\rho(t)$ -extendable path, its length is equal to  $|Z|$  by Lemma 12. Since  $\text{level}(Z') < \text{level}(\text{out}(t))$  holds,

$Z'[f, t]$  must terminate with  $\{p(t), t\}$  by Lemma 7(S1), and thus it is a  $\gamma(t)$ -alternating path by Lemma 7(S3). Then Lemma 7(S5) implies  $|Z'[f, t]| = \text{dist}^{\gamma(t)}(t)$  because if it does not hold one can obtain a  $\rho(t)$ -alternating path from  $f$  to  $z_t$  shorter than  $Z'$  by replacing the prefix  $Z'[f, t]$  of  $Z'$  with the shortest  $\gamma(t)$ -alternating path from  $f$  to  $t$ . Putting all the inequalities together, we obtain  $|R_{t, z_t}| \leq \text{dist}^{\gamma(z_t)}(z_t) - \text{dist}^{\gamma(t)}(t) \leq |Z'| - |Z'[f, t]| = |Z'[t, z_t]| = |Z|$ .

Next, we consider the second inequality. Let  $Y'$  be a shortest  $\rho(t)$ -alternating path from  $f$  to  $y_t$  of level less than  $\text{level}(\text{out}(t))$  by Lemma 3. Since it lies at the outside of  $T_t$ ,  $Y' \circ \{y_t, z_t\}$  is an alternating path from  $f$  to  $z_t$ . Then  $\text{dist}^{\gamma(z_t)}(z_t) \leq \text{dist}^{\rho(t)}(y_t) + 1$  holds. Since  $Y'[s, y_t]$  is a  $\rho(t)$ -extendable path, its length is equal to  $|Y|$ . Similarly, we obtain  $|Y'[f, s]| = \text{dist}^{\gamma(s)}(s)$ . Since  $s$  is an ancestor of  $t$ , we obtain  $\text{dist}^{\gamma(s)}(s) < \text{dist}^{\gamma(t)}(t)$  by Lemma 7(S4). Putting all the inequalities together, we have  $|Y| = |Y'[s, y_t]| = |Y'| - |Y'[f, s]| = \text{dist}^{\rho(t)}(y_t) - \text{dist}^{\gamma(s)}(s) > \text{dist}^{\gamma(z_t)}(z_t) - 1 - \text{dist}^{\gamma(t)}(t) \geq |R_{t, z_t}| - 1$ . That is,  $|Y| \geq |R_{t, z_t}|$  holds.  $\square$

The total running time is bounded as follows:

**Lemma 6** (Repeat). *Let  $g(h)$  be the worst-case round complexity of  $\text{EXTPATH}(s, t, \theta)$  over all  $(s, t, \theta)$  admitting a  $\theta$ -extendable path from  $s$  to  $t$  of length  $h$ . For any  $0 \leq h \leq \ell$ ,  $g(h) \leq 2h$  holds.*

*Proof.* We first see that any message transmitted in the algorithm does not suffer edge congestion, i.e., any two recursive invocations simultaneously executed do not simultaneously send two different messages through the same link with the same direction. A simple inductive argument shows that only vertices in  $V(T_s)$  sends a message in the run of  $\text{EXTPATH}(s, t, \theta)$ . Given two recursive invocations  $\text{EXTPATH}(s, t, \theta)$  and  $\text{EXTPATH}(s', t, \theta)$ ,  $V(T_s)$  and  $V(T_{s'})$  intersect if and only if  $s$  and  $s'$  have the ancestor-descendant relationship, which can occur only in the case of  $\text{EXTPATH}(s', t, \theta)$  is called at the inside of  $\text{EXTPATH}(s, t, \theta)$  or vice versa. That is, they never simultaneously run, and thus any two invocations simultaneously running do not send two messages from the same vertex.

We prove the time bound claimed in the lemma following the observation above. Specifically, we bound the running time of  $\text{EXTPATH}(s, t, \theta)$  by  $2h$  for any  $(s, t, \theta)$  admitting a  $\theta$ -extendable path from  $s$  to  $t$  of length  $h$ . The proof is based on the induction on  $h$ .

**(Basis)** The case of  $h = 0$ , the lemma obviously holds.

**(Inductive Step)** Suppose as the induction hypothesis that  $g(h') \leq 2h'$  holds for any  $0 \leq h' < h$ . We consider the following two cases:

**(Case 1)** For  $(s, t, \theta)$  such that  $\theta = \gamma(t)$  holds: By the correctness of  $\text{EXTPATH}$ ,  $(s, p(t))$  admits a  $\bar{\theta}$ -extendable path of length less than  $h$ . The induction hypothesis bounds the running time of the recursive call  $\text{EXTPATH}(s, p(t), \bar{\theta})$  by  $2(h - 1)$  rounds. Adding the cost of transmitting the triggering message (one round), we upper bound the running time of  $\text{EXTPATH}(s, t, \theta)$  by  $2h$ .

**(Case 2)** For  $(s, t, \theta)$  such that  $\theta = \overline{\gamma(v)}$  holds: By the correctness of  $\text{EXTPATH}$ ,  $(s, y_t)$  and  $(t, z_t)$  admit  $\rho(t)$ -extendable paths  $Y$  and  $Z$  respectively, which satisfy  $|Y| + |Z| + 1 = h$ . Then, by the induction hypothesis,  $\text{EXTPATH}(s, y_t, \rho(t))$  and  $\text{EXTPATH}(t, z_t, \rho(t))$  respectively run within  $2|Y|$  and  $2|Z|$  rounds. Since they run in parallel, the running time for the recursive calls is bounded by  $2 \max\{|Y|, |Z|\}$ . By Lemma 5, the length of  $R_{t, z_t} \circ \{y_t, z_t\}$ , which is the cost for transmitting the triggering message, is at most  $\min\{|Y| + 1, |Z| + 1\}$ . The total running time is at most  $2 \max\{|Y|, |Z|\} + \min\{|Y| + 1, |Z| + 1\} \leq 2(|Y| + |Z| + 1) = 2h$ .  $\square$

## 5 Concluding Remarks

In this paper, we presented a nearly linear-time exact algorithm for the maximum matching problem in the CONGEST model. This algorithm substantially improves the current best up-

per bound. To construct our algorithm, we proved a new structural lemma of alternating base trees (Lemma 3), which yields a very simple recursive construction of shortest augmenting paths in a fully-decentralized manner. We believe that this is a promising tool toward further improved algorithms.

Finally, we close this paper with a short remark on designing further improved (i.e., sublinear-time) algorithms. A major open problem following our result is to develop an algorithm of finding an augmenting path of length  $\ell$  within  $O(\ell^{1-\delta})$  rounds for a small constant  $\delta > 0$ . Since the iterative improvement of matchings can occur  $\Theta(n)$  times, a straightforward application of such an algorithm does not yield any sublinear-time algorithm. Fortunately, this issue is resolved by the  $(1-\epsilon)$ -approximate matching algorithm by Fischer, Slobodan, and Uitto running in  $\text{poly}(1/\epsilon, \log n)$  rounds: We first run it with  $\epsilon = n^{-\alpha}$  for a sufficiently small  $\alpha > 0$  such that its running time becomes sublinear. Starting from the output of the algorithm, we apply the sequential iterative improvement by augmenting paths. Due to the guarantee of  $(1-\epsilon)$ -approximation, the number of iterations is bounded by  $O(n^{1-\alpha})$ . If one can utilize an  $O(\ell^{1-\delta})$ -round augmenting path algorithm, the running time of whole iterative improvement process is sublinearly bounded. In summary, any  $O(\ell^{1-\delta})$ -round algorithm of finding an augmenting path of length  $\ell$  derives a sublinear-time CONGEST algorithm for the exact maximum matching problem.

## Acknowledgments

The first author was supported by JSPS KAKENHI Grant Numbers 23H04385, 22H03569, and 21H05854. The second author was supported by JSPS KAKENHI Grant Numbers 22K21277 and 23K16838. The third author was supported by JSPS KAKENHI Grant Numbers 20K19743 and 20H00605.

## References

- [1] Mohamad Ahmadi and Fabian Kuhn. Distributed maximum matching verification in CONGEST. In *34th International Symposium on Distributed Computing (DISC)*, pages 37:1–37:18, 2020.
- [2] Mohamad Ahmadi, Fabian Kuhn, and Rotem Oshman. Distributed approximate maximum matching in the CONGEST model. In *32nd International Symposium on Distributed Computing (DISC)*, pages 6:1–6:17, 2018.
- [3] Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of Algorithms*, 7(4):567–583, 1986.
- [4] Nir Bacrach, Keren Censor-Hillel, Michal Dory, Yuval Efron, Dean Leitersdorf, and Ami Paz. Hardness of distributed optimization. In *2019 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 238–247, 2019.
- [5] Reuven Bar-Yehuda, Keren Censor-Hillel, Mohsen Ghaffari, and Gregory Schwartzman. Distributed approximation of maximum independent set and maximum matching. In *36th annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 165–174, 2017.
- [6] Ran Ben-Basat, Ken-ichi Kawarabayashi, and Gregory Schwartzman. Parameterized distributed algorithms. In *33rd International Symposium on Distributed Computing (DISC)*, pages 6:1–6:16, 2018.

- [7] Claude Berge. Two theorems in graph theory. *Proceedings of the National Academy of Sciences of the United States of America*, 43(9):842–844, 1957.
- [8] Norbert Blum. A new approach to maximum matching in general graphs. In *International Colloquium on Automata, Languages, and Programming*, pages 586–597, 1990.
- [9] Yi-Jun Chang and Hsin-Hao Su. Narrowing the LOCAL-CONGEST gaps in sparse networks via expander decompositions. In *2022 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 301–312, 2022.
- [10] Jack Edmonds. Maximum matching and a polyhedron with 0,1-vertices. *Journal of Research of the National Bureau of Standards Section B Mathematics and Mathematical Physics*, page 125, 1965.
- [11] Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, pages 449–467, 1965.
- [12] Manuela Fischer, Slobodan Mitrović, and Jara Uitto. Deterministic  $(1 + \epsilon)$ -approximate maximum matching with  $\text{poly}(1/\epsilon)$  passes in the semi-streaming model and beyond. In *54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 248–260, 2022.
- [13] Harold N Gabow and Robert E Tarjan. Faster scaling algorithms for general graph matching problems. *Journal of the ACM (JACM)*, pages 815–853, 1991.
- [14] Mohsen Ghaffari, Fabian Kuhn, and Yannic Maus. On the complexity of local distributed graph problems. In *49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 784–797, 2017.
- [15] David G. Harris. Distributed local approximation algorithms for maximum matching in graphs and hypergraphs. In *60th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 700–724, 2019.
- [16] John E Hopcroft and Richard M Karp. An  $n^{5/2}$  algorithm for maximum matchings in bipartite graphs. *SIAM Journal on Computing*, pages 225–231, 1973.
- [17] Shang-En Huang and Hsin-Hao Su.  $(1 - \epsilon)$ -approximate maximum weighted matching in  $\text{poly}(1/\epsilon, \log n)$  time in the distributed and parallel settings. In *2023 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 44–54, 2023.
- [18] Amos Israeli and Alon Itai. A fast and simple randomized parallel algorithm for maximal matching. *Information Processing Letters*, pages 77–80, 1986.
- [19] Naoki Kitamura and Taisuke Izumi. A subquadratic-time distributed algorithm for exact maximum matching. *IEICE Transactions on Information and Systems*, 105(3):634–645, 2022.
- [20] Naoki Kitamura and Taisuke Izumi. Private communication, 2023.
- [21] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. The price of being near-sighted. In *17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 980–989, 2006.
- [22] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. Local computation: Lower and upper bounds. *Journal of the ACM (JACM)*, pages 1–44, 2016.

- [23] Zvi Lotker, Boaz Patt-Shamir, and Seth Pettie. Improved distributed approximate matching. *Journal of the ACM (JACM)*, pages 1–17, 2015.
- [24] Silvio Micali and Vijay V. Vazirani. An  $O(\sqrt{VE})$  algorithm for finding maximum matching in general graphs. In *21st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 17–27, 1980.
- [25] Vijay V. Vazirani. An improved definition of blossoms and a simpler proof of the MV matching algorithm. *CoRR*, abs/1210.4594, 2012.
- [26] Vijay V. Vazirani. A proof of the MV matching algorithm. *CoRR*, abs/2012.03582, 2020.

## A Proof Outline of Lemmas 1 and 2

### A.1 Lemma 1

Let  $A(M, \ell)$  be a CONGEST algorithm with the following property:

For any graph  $G = (V, E)$  and matching  $M \subseteq E$ ,  $A(M, \ell)$  finds a nonempty set of vertex-disjoint augmenting paths with high probability within  $O(\ell)$  rounds if  $(G, M)$  has an augmenting path of length at most  $\ell$ . Each vertex  $u \in V(G)$  outputs the predecessor and successor of the output augmenting path to which  $u$  belongs (if it exists).

The construction of the algorithm  $A$  is discussed in Section 3. The pseudocode of whole algorithm is presented in Algorithm 2. It basically follows the standard idea of centralized maximum matching algorithms, i.e., finding an augmenting path and improving the current matching iteratively. In the  $i$ -th iteration, the algorithm  $A(M, \ell)$  runs with  $\ell = \lfloor 2\hat{\mu}/(\hat{\mu} - i) \rfloor$ . This setting comes from Theorem 2. The improvement of the current matching by a given augmenting path is simply a local operation and is realized by flipping the labels of matching edges and non-matching edges on the path. The correctness and running time of Algorithm 2 are stated as follows:

**Lemma 14.** *Assume that there exists an algorithm  $A(M, \ell)$  specified above. Then Algorithm 2 constructs a maximum matching with high probability in  $\tilde{O}(\mu(G))$  rounds.*

*Proof.* The running time of the algorithm is bounded as follows (recall  $\hat{\mu} = O(\mu(G))$ ):

$$O\left(\sum_{i=0}^{\hat{\mu}-1} \left(\left\lfloor \frac{2\hat{\mu}}{\hat{\mu}-i} \right\rfloor\right)\right) = O\left(\hat{\mu} \sum_{i=1}^{\hat{\mu}} \frac{1}{i}\right) = O(\hat{\mu} \log \hat{\mu}) = \tilde{O}(\mu(G)).$$

We consider the correctness of the algorithm. Let  $s(i)$  be the matching size at the beginning of the  $i$ -th iteration. To prove the correctness, it suffices to show that  $s(\hat{\mu} - \mu(G) + i + 1) > s(\hat{\mu} - \mu(G) + i)$  for any  $i$  satisfying  $s(\hat{\mu} - \mu(G) + i) = i < \mu(G)$ . As  $\hat{\mu} \geq \mu(G)$ , this inductively implies that  $s(\hat{\mu} - \mu(G) + i) \geq i$  for any  $0 \leq i \leq \mu(G)$ , and hence  $s(\hat{\mu}) = \mu(G)$ .

Suppose that  $s(\hat{\mu} - \mu(G) + i) = i < \mu(G)$ . Then, in the  $(\hat{\mu} - \mu(G) + i)$ -th iteration, there exists an augmenting path of length at most  $\lfloor 2\mu(G)/(\mu(G) - i) \rfloor \leq \lfloor 2\hat{\mu}/(\mu(G) - i) \rfloor = \lfloor 2\hat{\mu}/(\hat{\mu} - (\hat{\mu} - \mu(G) + i)) \rfloor$  by Theorem 2. Thus the algorithm necessarily finds at least one augmenting path in this iteration, which implies  $s(\hat{\mu} - \mu(G) + i + 1) > s(\hat{\mu} - \mu(G) + i)$ .  $\square$



---

**Algorithm 2** Constructing a maximum matching in  $\tilde{O}(\mu(G))$  rounds.

---

- 1: **for**  $i = 0; i < \hat{\mu}; i++$  **do**
  - 2:   run the algorithm  $A(M, \ell)$  with  $\ell = \lfloor 2\hat{\mu}/(\hat{\mu} - i) \rfloor$  for  $O(\ell)$  rounds.
  - 3:   **if**  $A(M, \ell)$  finds a nonempty set of vertex-disjoint augmenting paths within  $O(\ell)$  rounds, **then**
  - 4:     improve the current matching using the set of vertex-disjoint augmenting paths.
- 

## A.2 Lemma 2

Although the original algorithm is designed for the verification of maximum matching, it provides each vertex with information on the length of alternating paths to the closest unmatched vertex. The following theorem holds.

**Theorem 3** (Ahmadi and Kuhn [1]). *Assume that a graph  $G = (V, E)$  and a matching  $M \subseteq E$  are given, and let  $W$  be the set of all unmatched vertices. There exist two  $O(\ell)$ -round randomized CONGEST algorithms  $MV(M, \ell, f)$  and  $PART(M, \ell)$  that output the following information at every vertex  $v \in V(G)$  with a probability of at least  $1 - 1/n^c$  for an arbitrarily large constant  $c > 1$ .*

1. *Given  $M$ , a nonnegative integer  $\ell$ , and a vertex  $f \in W$ ,  $MV(M, \ell, f)$  outputs the pair  $(\theta, \text{dist}_G^\theta(f, v))$  at each vertex  $v \in V(G)$  if  $\text{dist}_G^\theta(f, v) \leq \ell$  holds (if the condition is satisfied for both  $\theta = \text{odd}$  and  $\theta = \text{even}$ ,  $v$  outputs two pairs). The algorithm  $MV(M, \ell, f)$  is initiated only by the vertex  $f$  (with the value  $\ell$ ), and other vertices do not require information on the ID of  $f$  and value  $\ell$  at the beginning of the algorithm.*
2. *The algorithm  $PART(M, \ell)$  outputs a family of disjoint subsets  $V^1, V^2, \dots, V^N$  of  $V(G)$  (as the label  $i$  for each vertex in  $V^i$ ) such that (a) the subgraph  $G^i$  induced by  $V^i$  contains exactly two unmatched vertices  $f^i$  and  $g^i$  and an augmenting path between  $f^i$  and  $g^i$  of length at most  $\ell$ , (b) for every vertex  $v \in V^i$ , the length of the shortest alternating path from  $f^i$  to  $v$  is at most  $\ell$  (and hence the diameter of  $G^i$  is  $O(\ell)$ ), and (c) every vertex in  $V^i$  knows the ID of  $f^i$ .*

While the original paper [1] only presents a single algorithm returning the outputs of both  $MV$  and  $PART$ , in a slightly weaker form, it is very straightforward to extend the result to the algorithms above (see [19] for the details of this extension). It is easy to transform any algorithm  $A'$  of Lemma 2 to the algorithm  $A$  satisfying the specification of Lemma 1. The transformation just follows (1) running  $PART(M, \ell)$  and (2) executing the algorithms  $MV(M, \ell, f)$  and  $A'$  at each subgraph  $G^i$  independently.

## B Precomputing Phase of Distributed Implementation

We recall the precomputed information required by  $EXTPATH$ :

- The parent and children of  $u$  in  $T$ , and the value of  $\gamma(u)$ .
- The IDs of two endpoints of  $\text{out}(u)$  (if  $\text{out}(u)$  is not virtual).
- Let  $R_{v, z_v}$  be the shortest path from  $v$  to  $z_v$  in  $T$ . If  $u \in V(R_{v, z_v})$ ,  $u$  has the information for routing messages from  $v$  to  $z_v$ . More precisely,  $u$  knows the pair  $(z_v, u')$ , where  $u'$  is the successor of  $u$  in  $R_{v, z_v}$ .

The alternating base tree  $T$  is constructed within  $O(1)$  rounds by exchanging the information of  $\text{dist}^{\text{odd}}(\cdot)$  and  $\text{dist}^{\text{even}}(\cdot)$  between every neighboring vertex pair, because each vertex can determine its parent in  $T$  from the exchanged information. We refer to the height of  $T$  as  $\ell'$ . Obviously  $\ell' = O(\ell)$  holds.

The information of  $\text{out}(u)$  is obtained by the pipelined broadcast and aggregation over  $T$ : First, each node  $u$  constructs the list of all its ancestors. This task is implemented by the pipelined (downward) broadcast in  $T$ , which takes  $O(\ell)$  rounds. Then the constructed list is exchanged through every non-tree edge  $e = \{x, y\}$ . More precisely,  $x$  and  $y$  exchange the information of the ancestor list,  $\text{dist}^{\text{odd}}(\cdot)$ , and  $\text{dist}^{\text{even}}(\cdot)$ . Since the list size is bounded by the height of  $T$ , i.e.,  $O(\ell)$ , this information exchange also finishes within  $O(\ell)$  rounds. Then  $x$  and  $y$  can compute the level of  $e$ , and the ID of the least common ancestor of  $x$  and  $y$  (denoted by  $\text{lca}(e)$ ). We define  $F_i$  as the set of non-tree edges  $e$  such that the depth of  $\text{lca}(e)$  in  $T$  is  $i$ . Note that each node  $u$  can decide if  $e \in F_i \cap E^*(T_u)$  holds or not only from the information of  $\text{lca}(e)$  and its depth, because  $e \in F_i \cap E^*(T_u)$  holds if and only if the depth of  $\text{lca}(e)$  is  $i$  and  $\text{lca}(e)$  is an ancestor of  $u$ . For each  $0 \leq i \leq \ell' - 1$ , the algorithm informs each vertex  $u$  of the minimum-level edge in  $F_i \cap E^*(T_u)$  by a single-shot tree aggregation over  $T$ . The behavior of each vertex  $u$  in the  $i$ -th aggregation process is explained as follows:

- The vertex  $u$  is a leaf:  $u$  chooses the minimum-level edge in  $I(u) \cap (F_i \cap E^*(T_u)) = F_i \cap I(u)$ . If  $F_i \cap I(u)$  is empty, the dummy edge of level  $(\infty, \infty)$  is chosen. If two or more minimum-level edges exist, an arbitrary one of them is chosen.
- The vertex  $u$  is internal: Let  $X$  be the set of edges received from its children. The vertex  $u$  chooses the minimum-level edge in  $(X \cup I(u)) \cap (F_i \cap E^*(T_u))$ . If two or more minimum-level edges exist, the edge incident to  $u$  has a higher priority. For two edges with the same priority (i.e., two edges incident to  $u$ , or two edges not incident to  $u$ ), we use an arbitrary tie-breaking rule. Let  $\{v, w\}$  be the chosen edge, where  $w$  is the vertex in  $T_u$ . If  $w \neq u$  holds, the information of  $\{v, w\}$  comes from a child vertex  $u'$  of  $u$ , and thus  $u$  stores the routing information  $(w, u')$ .

By the standard pipeline technique, one can finish the tasks for all  $i$  in  $O(\ell)$  rounds. Finally, each vertex  $u$  computes its minimum outgoing edge as follows:

$$\begin{aligned} \text{out}(u) &= \text{minedge} \left( \bigcup_i F_i \cap E^*(T_u) \right) \\ &= \text{minedge} \left( \bigcup_i \text{minedge} (F_i \cap E^*(T_u)) \right), \end{aligned}$$

where  $\text{minedge}(F')$  is the minimum-level edge in  $F'$ , and the tie-breaking rule respects the rule of choosing canonical MOEs in Section 3.2.