On the Stability of Minimum-Weight Perfect Matching on the Line

Mark de Berg
Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands

Ulrike Schmidt-Kraepelin
Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands

Andree-Ovidiu Ștef
Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands

Abstract

Computing a minimum-weight perfect matching for a point set P in Euclidean space is a classic geometric optimization problem. We consider the problem in a dynamic setting, where pairs of points may be added to or removed from the set P. Our focus is on maintaining an approximately optimal solution without making too many changes to the solution. More precisely, we are interested in *k-stable algorithms*, which change at most k edges in the matching after each update to the set P. In other words, we consider an online setting (with insertions and deletions) with *bounded recourse*.

We study trade-offs between the stability of the algorithm and the approximation ratio of the maintained solution for point sets in $\mathbb{R}^1$. First, we present an $O(\sqrt{n})$-stable algorithm that maintains a 2-approximation, which we show to be optimal among all algorithms with sublinear stability. Second, we prove that any $o(\log n)$-stable algorithm has unbounded approximation ratio. Our lower bounds hold even in the insertion-only case, while our algorithm works in the fully dynamic case. Moreover, our lower bounds also hold for the bipartite variant of the problem.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Euclidean matching, stable approximation algorithms, dynamic algorithms, online algorithms, bounded recourse

Digital Object Identifier 10.4230/LIPIcs.ESA.2026.27

Related Version *Full Version*: `https://doi.org/10.48550/arXiv.2607.16137`

Funding *Mark de Berg*: MdB was supported by the Dutch Research Council (NWO) through Gravitation-grant NETWORKS-024.002.003.
Ulrike Schmidt-Kraepelin: USK was supported by the Dutch Research Council (NWO) under project number VI.Veni.232.254.

1 Introduction

Background. Computing minimum-weight perfect matchings is a classic problem in combinatorial optimization. In this problem, we are given an edge-weighted graph $\mathcal{G}$ on $2n$ vertices and the goal is to compute a *perfect matching* – a set of n disjoint edges – of minimum total weight. The problem, which we will simply refer to as the *matching problem*, can be solved in polynomial time using the well-known blossom algorithm by Edmonds [3, 4]. The currently fastest algorithm is due to Gabow [5] and runs in $O(n(m + n \log n))$ time, where m is the number of edges of $\mathcal{G}$. In the Euclidean variant of the problem, the input graph $\mathcal{G}$ is the complete graph on a set P of $2n$ points in Euclidean space and the edge weights correspond to the Euclidean distances between the points. Varadarajan [17], improving on work of Vaidya [16], showed that the matching problem in $\mathbb{R}^2$ can be solved in $O(n\sqrt{n}\log^5 n)$ time.

We are interested in the dynamic variant of the Euclidean matching problem, where pairs of points may be inserted into or deleted from the set P. One can, of course, compute an optimal solution from scratch after each update, but this is time consuming. More

importantly, the new solution may be completely different from the previous one. Since breaking up existing pairs in the matching and establishing new pairs may be expensive, switching to a completely new solution is undesirable. Thus, we want to maintain a perfect matching that does not change much after each update, while still having low weight.

This is similar to the goal of online algorithms, where the input is revealed over time – in other words, there are only insertions – and one has to maintain a feasible solution without revoking earlier decisions. For the matching problem as defined above, irrevocable decisions do not make much sense: the only option would be to always match the two newly arrived points to each other, which can give an arbitrarily bad approximation ratio. Hence, research on the online matching problem mainly focused on bipartite matching, where the point set P is partitioned into two subsets: a set S of servers known in advance, and a set C of clients that arrive over time and have to be matched to the servers. This problem is already quite difficult in $\mathbb{R}^1$, where it has been studied by various authors. The currently best algorithm for the online bipartite matching in $\mathbb{R}^1$ is by Raghvendra [15]; it has a competitive ratio of $O(\log n)$, improving an earlier $O(\log^2 n)$ bound by Raghvendra and Nayyar [12]. Recently, Peserico and Scquizzato [13] proved that any online algorithm for bipartite matching in $\mathbb{R}^1$ must have competitive ratio $\Omega(\sqrt{\log n})$, thus providing the first non-constant lower bound. There are also results on online bipartite matching in $\mathbb{R}^d$ for $d > 1$ and in arbitrary metric spaces; see for example [8, 9, 14]. Work that is closely related to our setting is by Gupta, Krishnaswamy, and Sandeep [7], who show how to maintain a 3-approximation for the bipartite matching problem in $\mathbb{R}^1$, while changing $O(\log n)$ edges in the matching upon each update. (They also obtain a result for general metrics, where the approximation ratio increases to $O(\log n)$, and for the case where clients as well as servers can arrive or depart; in the latter scenario, the number of changes becomes $O(\log \sigma)$, where σ is the spread of the metric.) A similar result was achieved by Megow and Nölke [11]. Recently, Goranci et al. [6] also study the closely related bipartite matching problem in any fixed-dimensional Euclidean space, where both clients and servers can arrive or depart. They show that, using randomization techniques and two quadtree-like data structures, it is possible to achieve an algorithm with an $O((1/\varepsilon)n^{\varepsilon})$ update runtime while maintaining an expected $O(1/\varepsilon)$-approximation for any $0 < \varepsilon \leqslant 1$. Moreover, they provide a lower bound of 2 on the approximability of an algorithm with a sublinear amortized update runtime, which is similar to one of the results presented in this paper.

Our interest, however, lies in non-bipartite Euclidean matching. Thus, we must allow some changes to the existing matching when new points are inserted or deleted. In other words, using terminology from online algorithms, we must allow *recourse*, similar to the papers by Gupta, Krishnaswamy, and Sandeep [7] and Megow and Nölke [11] we just mentioned. The question then becomes: how many changes are required to maintain a solution that has a good approximation ratio? Or, more generally: which trade-offs can we obtain between the number of changes we are allowed to make and the approximation ratio?

A related problem was studied by Matuschke, Schmidt-Kraepelin, and Verschae [10], who study non-bipartite matching in the following setting. Instead of considering a sequence of insertions of point pairs, they consider a scenario with two stages. In the first stage, a set P_1 of arbitrary (even) size is given, for which the algorithm must compute a perfect matching M_1. In the second stage, a set P_2 of $2n$ points is given, after which the algorithm must compute a perfect matching M_2 for $P_1 \cup P_2$. The goal is that the maximum approximation ratio of M_1 and M_2, denoted by α, is small, while the number of edges removed from M_1 in the second stage is small. For point sets in $\mathbb{R}^1$, they obtain approximation ratio $\alpha = 10$ with $|M_1 \setminus M_2| \leqslant 2n$. When n is known in advance, they improve the result to $\alpha = 3$ and $|M_1 \setminus M_2| \leqslant n$; this latter result actually holds in arbitrary metric spaces.

Our results. We consider the matching problem for a fully dynamic point set P, focusing on trade-offs between the number of changes to the matching after each update and the approximation ratio. Following previous work on the bipartite setting, we study this fundamental problem in $\mathbb{R}^1$, where it already appears to be quite challenging. We will state our results using the terminology introduced by De Berg, Sadhukhan, and Spieksma [1, 2]. To this end, consider a dynamic algorithm ALG that maintains a perfect matching on P. Let $P(t)$ be the point set at time t, where we assume $P(0) = \emptyset$, let $M_{\text{alg}}(t)$ be the matching computed by ALG on $P(t)$, and let $\text{ALG}(t) := \text{weight}(M_{\text{alg}}(t))$ be the weight of $M_{\text{alg}}(t)$. Furthermore, let $M_{\text{opt}}(t)$ be an optimal (that is, minimum-weight) perfect matching on $P(t)$, and let $\text{OPT}(t) := \text{weight}(M_{\text{opt}}(t))$.

We say that ALG is an $f(n)$-*stable algorithm* if, for any sequence of updates to the set P, we have $|M_{\text{alg}}(t+1) \,\Delta\, M_{\text{alg}}(t)| \leqslant f(n)$ for all $t \geqslant 0$, where $n := \frac{1}{2} \cdot \max(|P(t)|, |P(t+1)|)$ and Δ denotes the symmetric difference. We refer to $f(n)$ as the *stability* of ALG. The stability is essentially the same as the *recourse* used for online algorithms: if $|M_{\text{alg}}(t+1) \,\Delta\, M_{\text{alg}}(t)| = k$ then the number of edges deleted from the previous matching is equal to $(k-1)/2$. We say that ALG is an α-*approximation algorithm* if $\text{ALG}(t) \leqslant \alpha \cdot \text{OPT}(t)$ for all $t \geqslant 0$.

Our first result, presented in Section 2, is an $O(\sqrt{n})$-stable 2-approximation algorithm. We also prove that the approximation ratio 2 is best possible, not just for $O(\sqrt{n})$-stable algorithms but for any algorithm that is $o(n)$-stable. This result relies on a construction that is very similar to the one used by Goranci et al. [6]. Since our result explicitly lower bounds the per-update stability rather than the amortized runtime over a sequence of updates, we include it for completeness. Our second result, presented in Section 3, is that no algorithm with bounded approximation ratio can have stability $o(\log n)$. The latter lower bound uses an interesting construction, where the locations of the points are given by a generating function derived from the binary representation of the arrival times. We also show how to adapt both lower bounds to the bipartite variant of the problem, where in each event the pair of points being inserted or deleted contains one server and one client. In this setting, a feasible solution is a perfect matching such that each edge in the matching connects a server to a client.

Notation and terminology. It will be convenient to view the matching problem in purely geometric terms. Thus, we are given a dynamic point set P in $\mathbb{R}^1$, and we want to maintain a set of $|P|/2$ edges – segments connecting points from P – such that every point in P is an endpoint of exactly one edge. We refer to such a set of edges as a *matching* on P, omitting the adjective "perfect" for convenience. If the matching has fewer than $|P|/2$ edges, we use the term *partial matching*. We denote the optimal matching and the matching maintained by our algorithm ALG on a point set P by $M_{\text{opt}}(P)$ and $M_{\text{alg}}(P)$, respectively; thus $M_{\text{opt}}(t)$ and $M_{\text{alg}}(t)$ are shorthands for $M_{\text{opt}}(P(t))$ and $M_{\text{alg}}(P(t))$.

We denote the edge between two points $p_i, p_j \in P$ by p_ip_j and we denote its length by $|p_ip_j|$. Hence, the weight of a (possibly partial) matching M is $\text{weight}(M) := \sum_{p_ip_j \in M} |p_ip_j|$. We denote the total length of a set I of intervals by $\|I\|$, and for a (possibly partial) matching M we sometimes use $\|M\|$ as a shorthand for $\text{weight}(M)$. For a (possibly partial) matching M and an interval $[x_1, x_2]$, we define $M \cap [x_1, x_2] := \{e \cap [x_1, x_2] : e \in M\}$ to be the parts of the edges from M that lie in $[x_1, x_2]$. Note that $M \cap [x_1, x_2]$ need not be a (partial) matching; it is just a set of segments. We say that an edge p_ip_j *covers* an edge p_kp_ℓ if $p_kp_\ell \subseteq p_ip_j$.

When specifying a point set, we always list its points from left to right. Thus, if we write $P := \{p_1, p_2, \ldots, p_{2n}\}$ then $p_1 < p_2 < \cdots < p_{2n}$. For $1 \leqslant i < 2n$, we call p_i the *predecessor* of p_{i+1} and we call p_{i+1} the *successor* of p_i. The *parity* of a point p_i in a set P is the parity

of its index. Points of odd parity are called *odd points* and points of even parity are called *even points*. Note that the parity of a point may change when P is being updated. An edge between consecutive points is called an *elementary edge*. It is easily seen that $M_{\text{opt}}(P)$ is unique and consists of elementary edges, as stated in the following observation.

▶ **Observation 1.** *Let* $P := \{p_1, \ldots, p_{2n}\}$. *Then* $M_{\text{opt}}(P) = \{p_1p_2, p_3p_4, \ldots, p_{2n-1}p_{2n}\}$.

2 An $O(\sqrt{n})$-stable 2-approximation algorithm

Our algorithm maintains a decomposition $\mathcal{B}(t)$ of the current point set $P(t)$ into $O(\sqrt{n})$ *blocks* – sets of consecutive points in P – of size $O(\sqrt{n})$ each, and then handles the insertion (or deletion) of a pair q_1, q_2 as follows. Let $B(q_1)$ and $B(q_2)$ be the blocks containing q_1 and q_2, respectively. We will reconstruct the matchings in $B(q_1)$ and $B(q_2)$ from scratch and only make $O(1)$ changes to each of the other blocks. While this idea is simple, making it work is highly non-trivial. Next we describe the various invariants we need to maintain and prove that they guarantee a good approximation ratio. Then we prove the existence of a so-called *good pair of splitters* for a block, which determines the matching inside a block. Finally, we explain how to maintain the invariants with an $O(\sqrt{n})$-stable algorithm.

2.1 The structure

The block decomposition. Let $P := P(t)$ be the current point set. Recall that the points from P are numbered as $p_1, \ldots, p_{2n}$, from left to right. A *block* is defined as a subset of consecutive points from P, that is, a block is a subset $B \subset P$ of the form $\{p_i, \ldots, p_j\}$ for indices $1 \leqslant i \leqslant j \leqslant 2n$. We denote the leftmost point of a block B by $\ell(B)$ and the rightmost point by $r(B)$. With a slight abuse of notation, and when P is clear from the context, we will write $M_{\text{alg}}(B)$ to denote the edges in $M_{\text{alg}}(P)$ with both endpoints in the same block B. Similarly, we write $M_{\text{opt}}(B)$ for the edges in $M_{\text{opt}}(P)$ with both endpoints in B.

A *block decomposition* of P is an ordered set $\mathcal{B} := \{B_1, B_2, \ldots, B_{|\mathcal{B}|}\}$ of blocks that defines a partition of P, where B_1 is the leftmost block and $r(B_i)$ is the predecessor of $\ell(B_{i+1})$ for all $1 \leqslant i < |\mathcal{B}|$. The $O(\sqrt{n})$ stability parameter will be ensured by an invariant stating that for any block B in the block decomposition $\mathcal{B}(t)$ maintained by our algorithm, it holds that $m \leqslant |B| \leqslant 2m + 3$, where $m := \left\lfloor \sqrt{|P(t)|} \right\rfloor$. This can be maintained efficiently using two additional technical invariants; see the full paper for details.

Connectors between blocks. The matching M_{alg} maintained by our algorithm will mainly use intra-block edges, that is, edges that connect points within the same block. Our decomposition invariant does not guarantee that blocks are of even size, however, so we may also need some inter-block edges. We will only use inter-block edges connecting the rightmost point of a block B_i to the leftmost point of the next block B_{i+1}. We call such inter-block edges *connectors*. The next invariant specifies which connectors we use and shows that the set of connectors used by M_{alg} is uniquely determined by the current block decomposition.

> **Connector invariant.** Let B_i and B_{i+1} be consecutive blocks in the current block decomposition $\mathcal{B}$, and let $p_j := r(B_i)$. Then the matching M_{alg} satisfies:
> **(C)** $p_jp_{j+1} \in M_{\text{alg}}$ iff p_j is an odd point.

Note that B_{i+1} must exist when $r(B_i)$ is an odd point. Observation 1 and the connector invariant together imply that a connector p_jp_{j+1} is used in $M_{\text{alg}}(P)$ iff it is used in $M_{\text{opt}}(P)$.

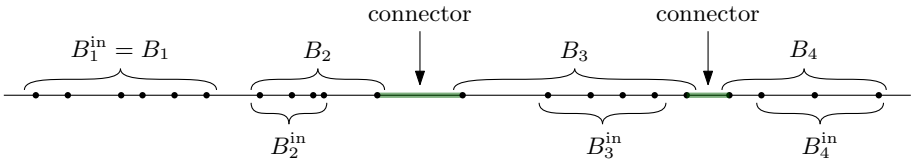

Figure 1 A block decomposition.

Unwrapped, inner-wrapped, and outer-wrapped blocks. We now come to the heart of our construction, which is how we match points within a block $B_i \in \mathcal{B}$. Let $p_j := \ell(B_i)$ and $p_k := r(B_i)$. Recall that the connector $p_{j-1}p_j$ is an edge in M_{alg} iff p_j is an even point. Similarly, p_kp_{k+1} is an edge in M_{alg} iff p_k is an odd point. This leads us to define $B_i^{\text{in}} := \{p_a, \ldots, p_b\}$, where $a = j$ if j is odd and $a = j + 1$ if j is even, and $b = k$ if k is even and $b = k - 1$ if k is odd. (If $|B_i| \leqslant 2$ then $B_i^{\text{in}} = \emptyset$.) Thus, B_i^{in} contains the points from B_i that still need to be matched after adding the connectors; see Figure 1. Note that $|B_i^{\text{in}}|$ must be even, as it starts at an odd point and ends at an even point. We call B_i^{in} the *interior* of B_i; we may have $B_i^{\text{in}} = B_i$. Note that B_i^{in} may change even if B_i does not change, because the insertion of a point in a block B_j with $j < i$ causes a parity change for $\ell(B_i)$.

We define three types of blocks, each using a specific matching in its interior. The type of any given block – in other words, which matching M_{alg} uses in the interior of the block – will depend on the parity of $\ell(B_i)$, as explained later. The three types are called unwrapped, inner-wrapped, and outer-wrapped. To define these types for a block $B_i \in \mathcal{B}$, we first define what it means for a sub-block $B \subseteq B_i$ to be wrapped or unwrapped; see Figure 2.

▶ **Definition 2.** *Let $B_{s,t} := \{p_s, \ldots, p_t\}$ be an even-sized sub-block of a block $B_i \in \mathcal{B}$. Define*

$$M_{\text{unw}}(B_{s,t}) := \{p_sp_{s+1}, p_{s+2}p_{s+3}, \ldots, p_{t-1}p_t\}$$

and

$$M_{\text{wr}}(B_{s,t}) := \{p_sp_t\} \cup \{p_{s+1}p_{s+2}, p_{s+3}p_{s+4}, \ldots, p_{t-2}p_{t-1}\}.$$

We call $B_{s,t}$ unwrapped *if $M_{\text{alg}}(B_{s,t}) = M_{\text{unw}}(B_{s,t})$ and* wrapped *if $M_{\text{alg}}(B_{s,t}) = M_{\text{wr}}(B_{s,t})$.*

Note that if a sub-block $B_{s,t}$ is unwrapped and p_s is an odd point, then $M_{\text{alg}}(B_{s,t}) = M_{\text{opt}}(B_{s,t})$ by Observation 1. The definition of an unwrapped block is straightforward.

- *Unwrapped blocks.* An unwrapped block B_i is a block such that B_i^{in} is unwrapped.

Inner- and outer-wrapped blocks will be defined with respect to two points $s_1(B_i)$ and $s_2(B_i)$ chosen from B_i^{in}, which we call the *splitters* of B_i. The choice of the splitters is crucial to obtain a good approximation ratio, and it will be discussed in detail below. For now, it suffices to know that splitters have the following property. Let $B_i^{\text{in}} := \{p_a, \ldots, p_b\}$ and let $s_1(B_i) := p_s$ and $s_2(B_i) := p_t$. Then $p_a \leqslant p_s < p_t \leqslant p_b$ and $\{p_s, \ldots, p_t\}$ is an even-sized sub-block. Inner- and outer-wrapped blocks are now defined as follows; see also Figure 3.

- *Inner-wrapped blocks.* An inner-wrapped block is a block $B_i \in \mathcal{B}$ where M_{alg} is such that $\{p_a, \ldots, p_{s-1}\}$ and $\{p_{t+1}, \ldots, p_b\}$ are unwrapped and $\{p_s, \ldots, p_t\}$ is wrapped.
- *Outer-wrapped blocks.* An outer-wrapped block is a block $B_i \in \mathcal{B}$ where M_{alg} is such that $\{p_a, \ldots, p_s\}$ and $\{p_t, \ldots, p_b\}$ are wrapped and $\{p_{s+1}, \ldots, p_{t-1}\}$ is unwrapped.

Figure 2 An unwrapped sub-block (left) and a wrapped sub-block (right).

Observe that $\{p_a, \ldots, p_{s-1}\}$ must have even size for an inner-wrapped block to be well-defined. Since p_a is an odd point by definition, $\{p_a, \ldots, p_{s-1}\}$ has even size iff $s_1(B_i)$ is an odd point. (Recall that $p_s = s_1(B_i)$.) Similarly, $s_1(B_i)$ must be an even point for an outer-wrapped block to be well-defined. This does not mean that the type of a block B_i is fixed as long as B_i and $s_1(B_i)$ do not change. Indeed, the parity of $s_1(B_i)$ can change due to an insertion into a block B_j for some $j < i$. When this happens, we must *toggle* B_i, that is, we must change its type from inner-wrapped to outer-wrapped, or vice versa. As illustrated in Figure 3, toggling a block only requires $O(1)$ changes in M_{alg}. (The figure illustrates the situation where $|B_i|$ is odd, but toggling also requires only $O(1)$ changes when $|B_i|$ is even.)

We want to pick $s_1(B_i)$ and $s_2(B_i)$ such that we have a good approximation ratio regardless if B_i is inner-wrapped or outer-wrapped. This leads to the following definition.

▶ **Definition 3.** *Let B_i be a block with $|B_i| \geqslant 6$ and define $\textsc{opt}(B_i^{\text{in}}) := \text{weight}(M_{\text{opt}}(B_i^{\text{in}}))$. Then two points $s_1(B_i) := p_s$ and $s_2(B_i) := p_t$ with $p_s < p_t$ form a* good pair of splitters *if*

(i) *The set $\{p_s, \ldots, p_t\}$ is an even-sized sub-block of B_i.*

(ii) *If p_s is an odd point in the current set P then the total weight of the inner-wrapped matching of B_i is at most $2 \cdot \textsc{opt}(B_i^{\text{in}})$, and if p_s is an even point in the current set P then the total weight of the outer-wrapped matching of B_i is at most $2 \cdot \textsc{opt}(B_i^{\text{in}})$.*

Our algorithm will maintain to the following invariant.

Wrapping invariant. Let B_i be a block in the current block decomposition $\mathcal{B}$. The matching M_{alg} satisfies the following property for a good pair of splitters $s_1(B_i), s_2(B_i)$.

(W.1) If $|B_i| < 6$ then B_i is unwrapped.

(W.2) If $|B_i| \geqslant 6$ and $s_1(B_i)$ is an odd point then B_i is inner-wrapped.

(W.3) If $|B_i| \geqslant 6$ and $s_1(B_i)$ is an even point then B_i is outer-wrapped.

In the next subsection we show that a good pair of splitters always exists, but first we show that the wrapping and connector invariants together imply that M_{alg} is a 2-approximation.

▶ **Lemma 4.** *Let $M_{\text{alg}}(P)$ be a matching adhering to the wrapping invariant and the connector invariant. Then $M_{\text{alg}}(P)$ is a perfect matching on P such that* $\text{weight}(M_{\text{alg}}(P)) \leqslant 2 \cdot \textsc{opt}(P)$.

2.2 Finding a good pair of splitters

Let B be a block of size at least 6 for which we want to find a good pair of splitters. It will be convenient to label the points in B from left to right as $b_0, b_1, \ldots, b_{z+1}$, but it is important to realize that there may be other points from P before B. Thus, the parity of b_0 can still be odd or even. Recall that the splitters should be chosen from B^{in}, the interior of

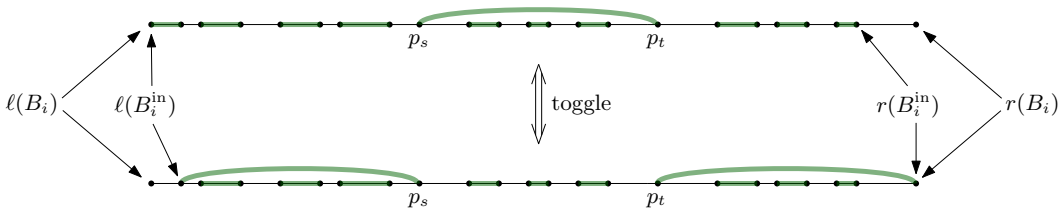

Figure 3 An inner-wrapped block (top) and an outer-wrapped block (bottom). If the parity of $\ell(B_i)$ – and, hence, of p_s – changes due to an insertion in an earlier block, then the type of B_i changes from inner-wrapped to outer-wrapped, or vice versa. We then say that B_i is toggled. Note that in the situation at the bottom, the point $\ell(B_i)$ must be incident to a connector.

the block B. Depending on the parity of b_0 (which can change), b_0 and b_{z+1} may or may not be part of B^{in}, so we have to choose the splitters from $\{b_1, \ldots, b_z\}$. To simplify the notation, we assume that $b_1 = 0$ and $b_z = 1$. This is without loss of generality, as scaling the instance does not change the approximation ratio we achieve inside the block.

Let $x_1, x_2 \in [0,1]$ be a potential pair of splitters, with $x_1 < x_2$. Then x_1 and x_2 split the interval $[0,1]$ into two parts: an *inner part* $[x_1, x_2]$ and an *outer part* $[0, x_1) \cup (x_2, 1]$. We first find a suitable partition into an inner and an outer part in a continuous setting, where we can pick the splitting points x_1 and x_2 anywhere in the interval $[0,1]$. After that, we explain how to obtain a good pair of splitters $b_s, b_t \in \{b_1, , \ldots, b_z\}$ from the pair x_1, x_2.

$$\text{Define} \quad I_{\text{even}} := [b_1, b_2] \cup [b_3, b_4] \cup \cdots \cup [b_{2y-1}, b_{2y}], \text{where } y = \lfloor z/2 \rfloor.$$

Thus, I_{even} contains the edges that are in $M_{\text{opt}} \cap [0,1]$ when b_0 is an even point. We now define a *target length* for either the inner or the outer part. Recall that $\|I_{\text{even}}\|$ denotes the total length of the edges in I_{even}; thus, $\|I_{\text{even}}\| = \text{weight}(I_{\text{even}})$. The target length will depend on $\|I_{\text{even}}\|$ and is defined as follows. Let $\phi : [0,1] \to [0,1]$ be given by

$$\phi(\ell) := \begin{cases} \frac{\ell}{2(1-\ell)} & \text{if } 0 \leqslant \ell \leqslant \frac{1}{2} \\ \frac{3\ell-1}{2\ell} & \text{if } \frac{1}{2} < \ell \leqslant 1 \end{cases}$$

The target length is now defined $\phi(\|I_{\text{even}}\|)$. To shorten the formulas we define $W := \|I_{\text{even}}\|$ so that the target length becomes $\phi(W)$. (Unfortunately, there is no clear intuition behind the definition of $\phi(W)$; it is simply the function such that the computations work out.)

▶ **Lemma 5.** *For any block B, there exists $x_1, x_2 \in [0,1]$ with $x_1 < x_2$ such that*
(i) $\|I_{\text{in}}\| = \phi(W)$ *and* $\|I_{\text{in}} \cap I_{\text{even}}\| = W \cdot \phi(W)$, *or*
(ii) $\|I_{\text{out}}\| = \phi(W)$ *and* $\|I_{\text{out}} \cap I_{\text{even}}\| = W \cdot \phi(W)$,
where $I_{\text{in}} := [x_1, x_2]$ and $I_{\text{out}} := [0, x_1) \cup (x_2, 1]$ are the inner and outer parts defined by x_1, x_2. Moreover, a pair x_1, x_2 with the required property can be computed in $O(|B|)$ time.

Proof. For $x \in [0,1]$ we define an interval, or pair of intervals, I_x as follows:

$$I_x := \begin{cases} [x, x + \phi(W)] & \text{if } x \leqslant 1 - \phi(W) \\ [0, x - 1 + \phi(W)] \cup [x, 1] & \text{if } 1 - \phi(W) < x \leqslant 1. \end{cases}$$

Note that $\|I_x\| = \phi(W)$ for all x. Intuitively, as x increases from 0 to 1, the interval I_x moves to the right until it hits the point 1, after which it "wraps around" and consists of two subintervals. Define $f(x) := \|I_x \cap I_{\text{even}}\|$ and consider the integral $\int_0^1 f(x)\,\mathrm{d}x$. Note that a point $y \in [0,1]$ is contained in $I_x \cap I_{\text{even}}$ iff x is contained in the set $I_{\text{even}} - \phi(W)$, that is, in the set obtained by shifting I_{even} backwards over a distance $\phi(W)$ and wrapping around at $x = 0$ where necessary. Thus, each $y \in I_{\text{even}}$ contributes a total value $\phi(W)$ to the integral. Hence,

$$\int_0^1 f(x)\,\mathrm{d}x = \|I_{\text{even}}\| \cdot \phi(W) = W \cdot \phi(W),$$

where the second equality follows from simply plugging in the definition of W. Thus, $\min_{0 \leqslant x \leqslant 1} f(x) \leqslant W \cdot \phi(W) \leqslant \max_{0 \leqslant x \leqslant 1} f(x)$, and since $f(x)$ is continuous, this implies that there exists an x^* such that $f(x^*) = W \cdot \phi(W)$. To satisfy the conditions of the lemma we can therefore set $x_1 := x^*$ and $x_2 := x^* + \phi(W)$ if $x^* \leqslant 1 - \phi(W)$, and we can set $x_1 := x^* - 1 + \phi(W)$ and $x_2 := x^*$ otherwise.

Observe that $f(x)$ is a piecewise linear function with $O(|B|)$ pieces, which we can construct in $O(|B|)$ time, after which we can find x^* in $O(|B|)$ time. ◀

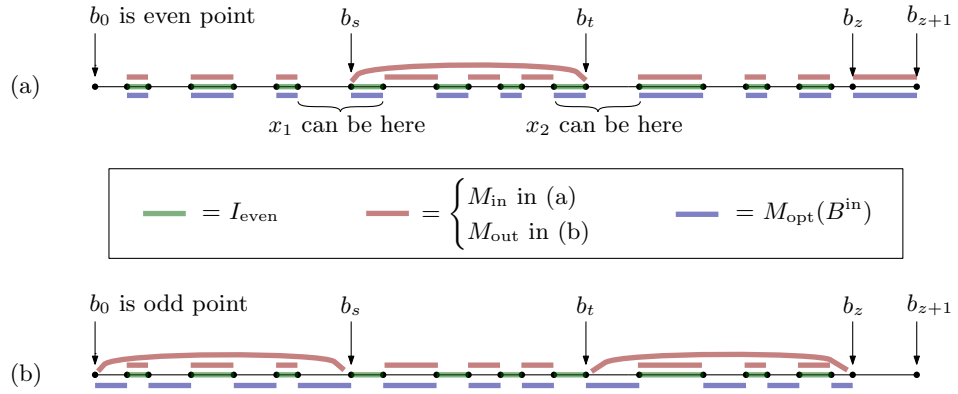

Figure 4 Illustration for the proof of Lemma 6.

Let x_1, x_2 be a pair of points with the property stated in Lemma 5 and let b_i and b_j such that $b_i \leqslant x_1 < b_{i+1}$ and $b_j < x_2 \leqslant b_{j+1}$. Note that $b_i, b_{i+1}, b_j, b_{j+1} \in \{b_1, \ldots, b_z\}$. In the full paper we show that $i \neq j$. We now choose the splitter $b_s \in \{b_i, b_{i+1}\}$ as follows.

(S.1) If x_1 and x_2 have property (i) from Lemma 5, we pick $b_s := b_i$ if $\{b_0, \ldots, b_i\}$ has even size and we pick $b_s := b_{i+1}$ if it has odd size. Thus, b_0 and b_s have opposite parity.

(S.2) If x_1 and x_2 have property (ii) from Lemma 5, we pick $b_s = b_i$ if $\{b_0, \ldots, b_i\}$ has odd size and we pick $b_s := b_{i+1}$ if it has even size. Thus, b_0 and b_s have the same parity.

After picking b_s, we pick $b_t \in \{b_j, b_{j+1}\}$ such that $\{b_s, \ldots, b_t\}$ has even size. The next lemma shows that b_s, b_t is a good pair of splitters when b_s and b_t are chosen according to **(S.1)**; in the full paper we show that the case where b_s and b_t are chosen according to **(S.2)** can be handled similarly.

▶ **Lemma 6.** *If b_s and b_t are chosen according to **(S.1)**, they are a good pair of splitters.*

Proof. The sub-block $\{b_s, \ldots, b_t\}$ has even size by construction, so it remains to prove property (ii) of Definition 3.

Case I: b_s is currently an odd point in P. Let M_{in} denote the inner-wrapped matching of B^{in}. We must show that $\text{weight}(M_{\text{in}}) \leqslant 2 \cdot \text{OPT}(B^{\text{in}})$. Since we picked b_s according to **(S.1)**, we know that the parity of b_0 is opposite to the parity of b_s. Hence, b_0 is an even point, as in Figure 4(a). Depending on whether $|B|$ is odd or even, B^{in} ends at b_{z+1} and we have $\text{OPT}(B^{\text{in}}) = \|I_{\text{even}}\| + |b_z b_{z+1}| = W + |b_z b_{z+1}|$, or B^{in} ends at b_z and we have $\text{OPT}(B^{\text{in}}) = \|I_{\text{even}}\| = W$. The former case is shown in Figure 4(a). For now, assume the latter case, that is, assume that $|B|$ is even so that $\text{OPT}(B^{\text{in}}) = W$. We can then bound $\text{weight}(M_{\text{in}})$ as follows, where (as before) $I_{\text{in}} := [x_1, x_2]$.

$$\text{weight}(M_{\text{in}}) \leqslant W + 2 \cdot (\|I_{\text{in}}\| - \|I_{\text{in}} \cap I_{\text{even}}\|) = W + 2 \cdot (\phi(W) - W \cdot \phi(W)),$$

where the last equality follows from property (i) in Lemma 5. We have two cases, depending on the value of W.

- If $W \leqslant \frac{1}{2}$ then $\phi(W) = \frac{W}{2(1-W)}$ and so

  $$2 \cdot (\phi(W) - W \cdot \phi(W)) = 2(1-W) \cdot \tfrac{W}{2(1-W)} = W.$$

  We can conclude that $\text{weight}(M_{\text{in}}) \leqslant 2W = 2 \cdot \text{OPT}(B^{\text{in}})$, as desired.
- On the other hand, if $W > \frac{1}{2}$ then $\phi(W) = \frac{3W-1}{2W}$ and so

  $$2 \cdot (\phi(W) - W \cdot \phi(W)) = 2(1-W) \cdot \tfrac{3W-1}{2W} = \tfrac{(1-W)(3W-1)}{W}$$

  It is easily verified that $\frac{(1-W)(3W-1)}{W} \leqslant W$ for any $0 < W \leqslant 1$. Hence, $\text{weight}(M_{\text{in}}) \leqslant 2W = 2 \cdot \text{OPT}(B^{\text{in}})$.

In the derivations above we assumed $|B|$ is even. Now suppose $|B|$ is odd. Then, compared to the case where $|B|$ is even, both $\text{weight}(M_{\text{in}})$ and $\text{OPT}(B^{\text{in}})$ increase by $|b_z b_{z+1}|$. This decreases the ratio $\text{weight}(M_{\text{in}})/\text{OPT}(B^{\text{in}})$, so we still have $\text{weight}(M_{\text{in}}) \leqslant 2 \cdot \text{OPT}(B^{\text{in}})$.

Case II: b_s is currently an even point in P. The proof is very similar to that of the previous case. Let M_{out} denote the outer-wrapped matching of B^{in}. We must show that $\text{weight}(M_{\text{out}}) \leqslant 2 \cdot \text{OPT}(B^{\text{in}})$. We now have that b_0 must be an odd point, as in Figure 4(b). As before, depending on whether $|B|$ is odd or even, B^{in} can end at b_z – this is the case shown in Figure 4(b) – or at b_{z+1}. We assume the former case; otherwise, $\text{weight}(M_{\text{in}})$ and $\text{OPT}(B^{\text{in}})$ both increase by $|b_z b_{z+1}|$, which only decreases the ratio $\text{weight}(M_{\text{in}})/\text{OPT}(B^{\text{in}})$.

If B^{in} ends at b_z then $\text{OPT}(B^{\text{in}}) = |b_0 b_1| \ + \ 1 - \|I_{\text{even}}\| = |b_0 b_1| \ + \ 1 - W$. We can now bound $\text{weight}(M_{\text{out}})$ as follows, where $I_{\text{out}} := [0, x_1) \cup (x_2, 1]$.

$$\begin{aligned}
\text{weight}(M_{\text{out}}) &\leqslant |b_0 b_1| \ + \ (1 - W) \ + \ 2 \cdot \|I_{\text{out}} \cap I_{\text{even}}\| \\
&= |b_0 b_1| + (1 - W) + 2 \cdot (\|I_{\text{even}}\| - \|I_{\text{in}} \cap I_{\text{even}}\|) \\
&= |b_0 b_1| \ + \ (1 - W) \ + \ 2 \cdot (W - W \cdot \phi(W)) \\
&= |b_0 b_1| \ + \ (1 - W) \ + \ \tfrac{2W(1-\phi(W))}{1-W} \cdot (1 - W)
\end{aligned}$$

Hence, it suffices to show that $\frac{2W(1-\phi(W))}{1-W} \leqslant 1$.

- If $W \leqslant \frac{1}{2}$ then $\phi(W) = \frac{W}{2(1-W)}$ and so

$$\tfrac{2W(1-\phi(W))}{1-W} = \tfrac{2W}{1-W}\left(1 - \tfrac{W}{2(1-W)}\right) = \tfrac{W(2-3W)}{(1-W)^2} \leqslant 1.$$

- On the other hand, if $W > \frac{1}{2}$ then $\phi(W) = \frac{3W-1}{2W}$ and so

$$\tfrac{2W(1-\phi(W))}{1-W} = \tfrac{2W}{1-W}\left(1 - \tfrac{3W-1}{2W}\right) = 1.$$

We conclude that the pair b_s, b_t has the required property in all cases. ◀

The next lemma summarizes the main result of this subsection.

▶ **Lemma 7.** *Let B be a block with $|B| \geqslant 6$. Then a good pair of splitters for B exists and can be computed in $O(|B|)$ time.*

2.3 Maintaining the structure

The idea behind the update algorithm is simple: we update the blocks containing the two inserted or deleted points from scratch and toggle the other blocks. Since the size of a block is $O(\sqrt{n})$ and the number of blocks is $O(\sqrt{n})$, and toggling a block changes only $O(1)$ edges, this means that the update algorithm is $O(\sqrt{n})$-stable. We show in the full paper how such a balanced block decomposition can be maintained, thus proving the following theorem.

▶ **Theorem 8.** *There is an $O(\sqrt{n})$-stable 2-approximation algorithm for the dynamic minimum-weight perfect matching problem in $\mathbb{R}^1$. The running time of the algorithm is $O(\sqrt{n})$ per insertion or deletion of a point pair.*

2.4 A lower bound on the approximation ratio

In the following, we show that our $O(\sqrt{n})$-stable 2-approximation algorithm is optimal in terms of its approximation ratio among all algorithms with sublinear stability.

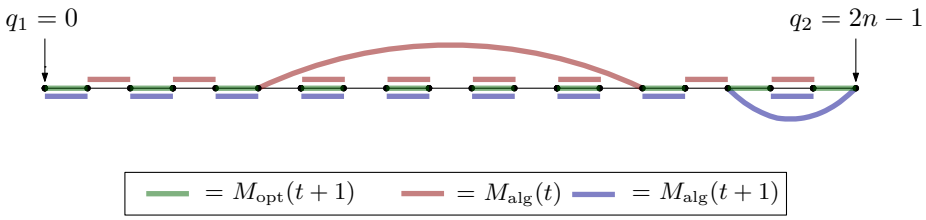

Figure 5 Counterexample used in Theorem 9.

▶ **Theorem 9.** *For no $\alpha < 2$, there is an algorithm for the dynamic minimum-weight perfect matching problem in $\mathbb{R}^1$ that is $o(n)$-stable and α-approximate.*

Proof. Let $\alpha < 2$ and assume for contradiction the existence of an α-approximate algorithm ALG that is $f(n)$-stable for some $f(n) = o(n)$. Choose $n \in \mathbb{N}$ large enough such that $f(n) \leqslant (1 - \frac{\alpha}{2})n$, and suppose that at time $t = n - 1$ we have $P(t) = \{1, 2, \ldots, 2(n-1)\}$ and that $P(t+1) = \{1, 2, \ldots, 2(n-1)\} \cup \{q_1, q_2\}$, where $q_1 := 0$ and $q_2 := 2n - 1$; see Figure 5. From the stability of ALG and because any edge in $P(t+1)$ has weight at least 1, we have

$$\begin{aligned} \text{weight}(M_{\text{alg}}(t) \cap M_{\text{alg}}(t+1)) &\geqslant |M_{\text{alg}}(t) \cap M_{\text{alg}}(t+1)| \\ &= |M_{\text{alg}}(t) \cup M_{\text{alg}}(t+1)| - |M_{\text{alg}}(t) \,\Delta\, M_{\text{alg}}(t+1)| \\ &\geqslant n - f(n) \\ &\geqslant \tfrac{\alpha}{2} n. \end{aligned}$$

Since $M_{\text{alg}}(t+1)$ matches all points in $P(t+1)$ while $M_{\text{alg}}(t)$ matches all points except $\{q_1, q_2\}$, the graph induced by the set $M_{\text{alg}}(t) \,\Delta\, M_{\text{alg}}(t+1)$ contains exactly two nodes with odd degree, namely q_1 and q_2. Thus, there exists a path $\pi \subseteq M_{\text{alg}}(t) \,\Delta\, M_{\text{alg}}(t+1)$ connecting q_1 and q_2. Clearly, $\text{weight}(\pi) \geqslant 2n - 1$. Hence,

$$\begin{aligned} \text{ALG}(t) + \text{ALG}(t+1) &= \text{weight}(M_{\text{alg}}(t) \Delta M_{\text{alg}}(t+1)) + 2 \cdot \text{weight}(M_{\text{alg}}(t) \cap M_{\text{alg}}(t+1)) \\ &\geqslant 2n - 1 + \tfrac{\alpha}{2}(2n-1) \\ &= \left(1 + \tfrac{\alpha}{2}\right)(2n-1) \\ &> \alpha \cdot (\text{OPT}(t) + \text{OPT}(t+1)), \end{aligned}$$

thus $\text{ALG}(t) > \alpha \cdot \text{OPT}(t)$ or $\text{ALG}(t+1) > \alpha \cdot \text{OPT}(t+1)$, giving the desired contradiction. ◀

Remark.. In the bipartite setting we can use the same construction, with points at even coordinates being servers and points at odd coordinates being clients. The algorithm is strictly more restricted in terms of the matchings it may maintain, whereas the weight of the optimal solution stays the same. Therefore Theorem 9 also holds in the bipartite setting.

3 A lower bound on the stability

In Subsection 2.4 we showed that any algorithm with a non-trivial upper bound on the stability must have an approximation ratio of at least 2. We now ask the reverse question: Is there a lower bound on the stability of any algorithm that has a bounded approximation

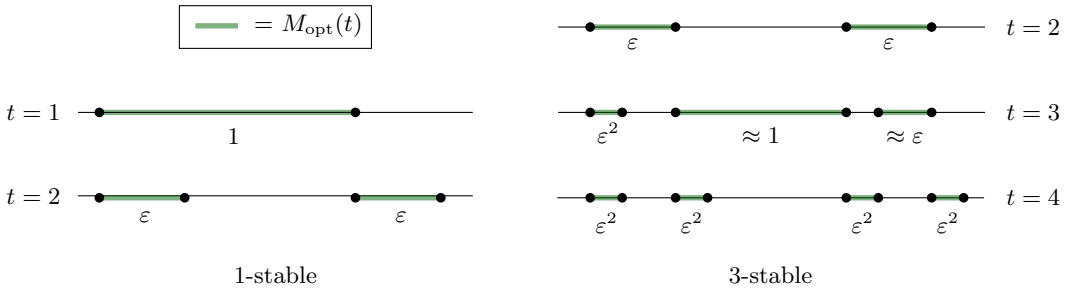

Figure 6 The counterexample for the 1-stable algorithm is $P(1) = \{0, 1\}$ with additional arrivals at $\{\varepsilon, 1 + \varepsilon\}$. For the 3-stable algorithm, the example starts with $P(2) = \{0, \varepsilon, 1, 1 + \varepsilon\}$ and adds points at $\{\varepsilon^2, 1 + \varepsilon^2\}$ at time $t = 3$ and $\{\varepsilon + \varepsilon^2, 1 + \varepsilon + \varepsilon^2\}$ at time $t = 4$.

guarantee? We call the approximation guarantee of an algorithm *bounded* if there exists some $\alpha : \mathbb{N} \to \mathbb{R}_{\geqslant 1}$ such that the algorithm is an $\alpha(n)$-approximation, that is, such that at any time t we have $\text{ALG}(t) \leqslant \alpha(n) \cdot \text{OPT}(t)$, where $n := \frac{1}{2}|P(t)|$.

A simple example shows that the unique 1-stable algorithm has unbounded approximation. Assume for contradiction that it is an $\alpha(n)$-approximation for some function α. Consider the point set $P := \{0, 1\}$ and the newly arriving points $q_1 := \varepsilon$ and $q_2 := 1 + \varepsilon$ for some $\varepsilon < \frac{1}{\alpha(2)}$. The optimal matching has weight 2ε but by 1-stability, ALG has to return the matching $\{(0, 1), (\varepsilon, 1 + \varepsilon)\}$, which has total weight $2 > \alpha(2) \cdot 2\varepsilon$.

The idea can be generalized to algorithms with higher stability. Doing so requires more than two events, however, since there exists a 3-approximate and 3-stable algorithm for the two-stage version of our problem [10]. Suppose ALG is 3-stable and $\alpha(n)$-approximate, and consider the scenario in Figure 6, where ε is chosen sufficiently small. At time $t = 2$ we have $\text{OPT} = 2\varepsilon$, forcing ALG to select both edges of weight ε. At time $t = 3$, the parity of the points requires ALG to select at least one edge from the leftmost three points to the rightmost three points. Finally, at time $t = 4$, ALG can only select edges of weight ε^2. Consequently, between time $t = 2$ and time $t = 4$, ALG must delete at least three edges and insert at least four edges. Thus, the total number of modifications is at least seven, whereas a 3-stable algorithm is permitted to make at most six changes in total. In the remainder of this section we generalize this idea to obtain a counterexample for any algorithm with an $o(\log n)$ stability, leading to the following result.

▶ **Theorem 10.** *Let ALG be an $o(\log n)$-stable algorithm for the dynamic minimum-weight perfect matching problem in $\mathbb{R}^1$. Then the approximation ratio of ALG cannot be bounded as a function of n.*

Let ALG be an algorithm that is an $\alpha(n)$-approximation for some function $\alpha : \mathbb{N} \to \mathbb{R}$. Assume for contradiction that ALG is also $f(n)$-stable for some function $f(n) = o(\log n)$. We can assume without loss of generality that $f(n)$ is non-decreasing[1] and there exists some constant n_0 that is a power of 2 and such that $f(n) < \log_2 n$ for all $n \geqslant n_0$. We present a sequence of n_0 events such that the total number of changes that ALG makes when processing this sequence is at least $n_0 \log n_0$, which contradicts that ALG is $f(n)$-stable as $\sum_{t=1}^{n_0} f(t) < n_0 \log_2 n_0$.

[1] If ALG is $f(n)$-stable then it is also $g(n)$-stable for $g(n) = \max_{i \leqslant n} f(i)$, which is non-decreasing.

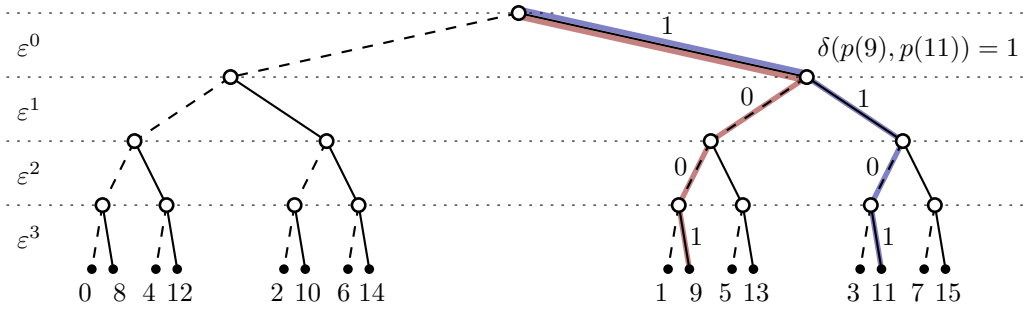

Figure 7 Trie interpretation for $P(8)$. Leaf labels correspond to times of insertion (e.g., at $t = 1$ indices $\{0, 1\} = \{2t - 2, 2t - 1\}$ are inserted). Leaves map to their position on the line, e.g., the leaf 9 is located at $p(9) = \varepsilon^0 + \varepsilon^3$.

The construction. For a positive integer k, define $[k] := \{1, \ldots, k\}$. Let $\alpha^* := \max_{t \in [n_0]} \alpha(t)$ and $\varepsilon := \frac{1}{2\alpha^* n_0}$. Given $k \in \mathbb{N}_0$, we define $k(i) \in \{0, 1\}$ to be the i-th least significant bit in the binary representation of k. So $k = \sum_{i=0}^{\mathrm{msbp}(k)} k(i) \cdot 2^i$, where $\mathrm{msbp}(k) := \lfloor \log_2 k \rfloor$ is the most significant bit position for $k > 0$ and $\mathrm{msbp}(0) := 0$. The least significant bit position $\mathrm{lsbp}(k)$ of k is the smallest i such that $k(i) = 1$. Moreover, let $p(k) := \sum_{i=0}^{\mathrm{msbp}(k)} k(i) \cdot \varepsilon^i$. The problem instance consists of n_0 insertions of pairs of points, where at time $t \in [n_0]$ we insert the points $p(2t - 2)$ and $p(2t - 1)$. Thus, for any t it holds that $P(t) = \{p(k) \mid k \in \{0, \ldots, 2t - 1\}\}$.

Trie interpretation. To gain intuition for our construction, it can be insightful to interpret the set $P(t)$ as the leaves of a binary trie, which we illustrate for $t = 8$ in Figure 7. Each level in the trie corresponds to a bit position in the binary representation of numbers in $\{0, \ldots, 2t-1\}$. Following the left edge of a node indicates the corresponding bit is 0, following the right indicates the bit is 1. Given any $k < 2t - 1$, we identify the corresponding leaf by following the bits in the binary representation of i, from least to most significant. For example, $11 = (1011)_2$, so we take the path 1101 to reach $p(11)$. The corresponding point on the line is located at $p(11) = \varepsilon^0 + \varepsilon^1 + \varepsilon^3$. It is important to note that, perhaps counterintuitively, the least significant bit in the binary representation corresponds to the coefficient of the most significant exponent, namely ε^0. The left-to-right ordering of the leaves in the trie corresponds to the ordering of the points in $P(t)$ on the line.

We define the *depth* of an edge $e := p(k)p(\ell)$ to be $\delta(e) := \mathrm{lsbp}(k \otimes \ell)$, where $\otimes$ is the bitwise XOR-operator. (Note that e is an edge between two of the points, not an edge in the trie.) In our trie representation, $\delta(e)$ corresponds to the level of the lowest common ancestor of the two leaves $p(k)$ and $p(\ell)$. We define E^d to be the set of all edges of depth d. The following lemma relates the depth of an edge to its weight.

▶ **Lemma 11.** *For any edge e it holds that* $\mathrm{weight}(e) \in \left(\frac{1}{2}\varepsilon^{\delta(e)}, \frac{3}{2}\varepsilon^{\delta(e)}\right)$.

Proof. Let $S := \sum_{i=1}^{\log_2(n_0)} \varepsilon^i$. By definition we have $\varepsilon < \frac{1}{3}$, so by the geometric-series formula we have $S < \frac{\varepsilon}{1-\varepsilon} < \frac{1}{2}$. Recall that $\delta(e)$ corresponds to the smallest exponent for which the two polynomials corresponding to $p(k)$ and $p(\ell)$ differ. Assume wlog that $k(\delta(e)) = 1$ and $\ell(\delta(e)) = 0$. Then

$$p(k) - p(\ell) \geqslant \varepsilon^{\delta(e)} - \sum_{i=\delta(e)+1}^{\log_2(n_0)} \varepsilon^i \geqslant \epsilon^{\delta(e)}(1 - S)$$

and

$$p(k) - p(\ell) \leqslant \sum_{i=\delta(e)}^{\log_2(n_0)} \varepsilon^i \leqslant \varepsilon^{\delta(e)}(1+S).$$

Hence,

$$|p(k) - p(\ell)| \in \left[\varepsilon^{\delta(e)}(1-S), \varepsilon^{\delta(e)}(1+S)\right] \subset \left(\tfrac{1}{2}\varepsilon^{\delta(e)}, \tfrac{3}{2}\varepsilon^{\delta(e)}\right),$$

which proves the claim. ◀

Note that Lemma 11, together with the fact that $\varepsilon < \frac{1}{3}$, implies that any edge of depth d is longer than any edge of depth $d+1$. Besides the depth of an edge, we introduce the notion of the *depth at time t*, defined as $\delta(t) := \mathrm{lsbp}(t)$. We show that there is a close link between $\delta(t)$ and the matchings $M_{\mathrm{opt}}(t)$ and $M_{\mathrm{alg}}(t)$. First, we give an upper bound on $\textsc{opt}(t)$.

▶ **Lemma 12.** *For any $t \in [n_0]$ it holds that $\textsc{opt}(t) \leqslant \varepsilon^{\delta(t)}t$.*

Proof. We prove the lemma by constructing a perfect matching M of weight $\varepsilon^{\delta(t)}t$. Let $m := 2^{\delta(t)}$. Define $M := \{p(k)p(k \otimes m) \mid 0 \leqslant k < 2t\}$. We begin by proving that M is a perfect matching of $P(t)$. By definition, we have that $2m$ divides $2t$. Since $2m = 2^{\delta(t)+1}$, for any integers $k < 2t$ and $\ell \geqslant 2t$, there exists $i > \delta(t)$ such that $k(i) \neq \ell(i)$. Additionally, $m(i) = 0$ for any $i > \delta(t)$, so we can define the function $f : \{0, \dots, 2t-1\} \to \{0, \dots, 2t-1\}$ with $f(k) = k \otimes m$. For any integers k and ℓ with $k \neq \ell$, there exists i such that $k(i) \neq \ell(i)$, but then also $f(k)(i) \neq f(\ell)(i)$ by definition of f, so f is injective. From the property that $f(f(k)) = (k \otimes m) \otimes m = k$ it also follows that f is surjective, hence bijective, and that $f = f^{-1}$. As such, M is a perfect matching of $P(t)$.

Let $p(k)p(\ell) \in M$ with $k < \ell$. Since $m = 2^{\delta(t)}$, we have that $\delta(t)$ is the only natural number for which $k(\delta(t)) \neq \ell(\delta(t))$. From this we first deduce that $\mathrm{weight}(p(k)p(\ell)) = \varepsilon^{\delta(t)}$, and therefore $\mathrm{weight}(M) = \varepsilon^{\delta(t)}t$. ◀

In addition, $M_{\mathrm{alg}}(t)$ must contain at least $2^{\delta(t)}$ edges of depth $\delta(t)$, and no edges at a lower depth. This builds on the following crucial observation, also illustrated in Figure 8.

▶ **Lemma 13.** *At any time $t \in [n_0]$, the point set $P(t)$ can be partitioned into $m := 2^{\delta(t)+1}$ sets $P_0, \dots, P_{m-1}$ such that:*

(i) *$|P_i|$ is odd for all $i \in \{0, \dots, m-1\}$,*

(ii) *$\delta(e) \leqslant \delta(t)$ for all edges e with endpoints in two different sets of the partition, and*

(iii) *$\delta(e) > \delta(t)$ for all edges e with endpoints in the same set of the partition.*

Proof. We partition the set $P(t)$ into subsets $P_0, \dots, P_{m-1}$, where we define

$$P_j := \{(p(k) \mid k \in \{0, \dots, 2t-1\} \text{ and } k \equiv j \pmod m\}.$$

We start by proving statement (i). Note that $\frac{2t}{m}$ is an odd integer since $m = 2^{\delta(t)+1} = 2^{\delta(2t)}$ is the largest power of two that divides $2t$. Moreover, for any $j \in \{0, \dots, m-1\}$ it holds that $|P_j| = \left|\{qm + j \mid q \in \{0, \dots, \frac{2t}{m} - 1\}\}\right| = \frac{2t}{m}$, which is thus odd.

We continue with proving statement (ii). Let $e := p(k)p(\ell)$ with $k, \ell < 2t$ such that $p(k)$ and $p(\ell)$ are not included in the same set of the partition $P_0, \dots, P_{m-1}$. Thus, $k \not\equiv \ell \pmod m$ and by $m = 2^{\delta(t)+1}$ there exists $i \leqslant \delta(t)$ such that $k(i) \neq \ell(i)$. Thus, $\delta(e) = \mathrm{lsbp}(k \otimes \ell) \leqslant \delta(t)$.

Finally, we prove statement (iii). Let $e := p(k)p(\ell)$ with $k, \ell < 2t$ such that $p(k)$ and $p(\ell)$ are included in the same set of the partition $P_0, \dots, P_{m-1}$. Then, $k \equiv \ell \pmod m$ and by $m = 2^{\delta(t)+1}$ it holds that $k(i) = \ell(i)$ for all $i \leqslant \delta(t)$ and thus $\delta(e) = \mathrm{lsbp}(k \otimes \ell) > \delta(t)$. ◀

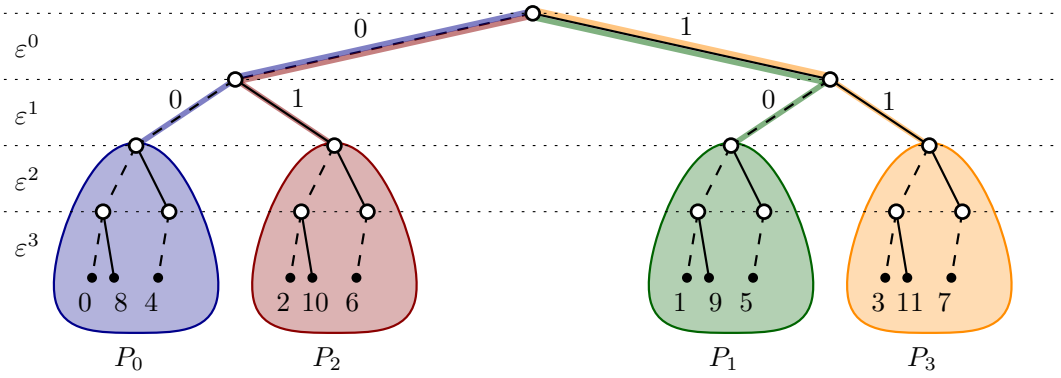

Figure 8 Illustration of Lemma 13. The figure depicts the partition of $P(6)$ into $m = 2^{\delta(6)+1} = 4$ subsets, each containing an odd number of points. Edges connecting two parts of the partition have depth at most 1.

For any depth $d \in \{0, \dots, \log_2 n_0\}$ and $t \in \{0, \dots, n_0\}$, let $M^d_{\text{alg}}(t) := M_{\text{alg}}(t) \cap E^d$ be the set of edges of $M_{\text{alg}}(t)$ at depth d. We can now provide a lower bound for $|M^{\delta(t)}_{\text{alg}}(t)|$.

▶ **Lemma 14.** *At every time t we have $|M^{\delta(t)}_{\text{alg}}(t)| \geqslant 2^{\delta(t)}$ and $|M^d_{\text{alg}}(t)| = 0$ for all $d < \delta(t)$.*

Proof. By Lemma 13(i), any feasible matching at time t contains at least $m/2 = 2^{\delta(t)}$ edges between different sets P_j, and by Lemma 13(ii) these edges have depth at most $\delta(t)$. Now assume for contradiction that $M_{\text{alg}}(t)$ contains an edge e with $\delta(e) < \delta(t)$. Then

$$\begin{aligned}
\text{ALG}(t) &\geqslant \text{weight}(e) \\
&> \tfrac{1}{2}\varepsilon^{\delta(t)-1} && \text{(by Lemma 11)} \\
&= \tfrac{1}{2} \cdot (2\alpha^* n_0) \cdot \varepsilon^{\delta(t)} && \text{(definition of } \varepsilon) \\
&\geqslant \alpha(t) \cdot \varepsilon^{\delta(t)} t && \text{(since } \alpha^* = \max_{t \in [n_0]} \alpha(t) \text{ and } t \leqslant n_0) \\
&\geqslant \alpha(t) \cdot \text{OPT}(t), && \text{(by Lemma 12)}
\end{aligned}$$

which contradicts that $M_{\text{alg}}(t)$ is an $\alpha(t)$-approximation. Hence, $|M^d_{\text{alg}}(t)| = 0$ for all $d < \delta(t)$ and $|M^{\delta(t)}_{\text{alg}}(t)| \geqslant 2^{\delta(t)}$. ◀

Lemma 14 allows us to lower bound the stability of ALG by summing over all $|M^{\delta(t)}_{\text{alg}}(t)|$ for all times $t \leqslant n_0$. We can do so because for any two times t_1 and t_2 with $\delta(t_1) = \delta(t_2)$ there exists a time $t_3 \in (t_1, t_2)$ such that $\delta(t_3) > \delta(t_1)$. Thus, at this point in time, ALG must have deleted all edges in $M^{\delta(t_1)}_{\text{alg}}(t_1)$ before including the edges in $M^{\delta(t_2)}_{\text{alg}}(t_2)$. The next lemma makes this precise.

▶ **Lemma 15.** *It holds that $\sum_{t=0}^{n_0-1} |M_{\text{alg}}(t) \Delta M_{\text{alg}}(t+1)| \geqslant n_0 \log_2 n_0$.*

Proof. We will argue that for any $d \in \{0, \dots, \log_2 n_0\}$, ALG must insert or delete at least n_0 edges of depth d. To this end, let $C_d := \frac{n_0}{2^d}$. Observe that time $t \leqslant n_0$ is of depth d if and only if $t = j \cdot 2^d$ for some odd $j \in \{0, \dots, C_d\}$. Moreover, when j is even, then $\delta(t) > d$ for $t = j \cdot 2^d$. Therefore, by Lemma 14 it holds that $|M^d_{\text{alg}}(j \cdot 2^d) \Delta M^d_{\text{alg}}((j+1)2^d)| \geqslant 2^d$ for any $j \in \{0, \dots, C_d\}$. We get

$$\sum_{t=0}^{n_0-1} |M^d_{\text{alg}}(t) \Delta M^d_{\text{alg}}(t+1)| \geqslant \sum_{j=0}^{C_d-1} |M^d_{\text{alg}}(j \cdot 2^d) \Delta M^d_{\text{alg}}((j+1)2^d)| \geqslant C_d \cdot 2^d = n_0.$$

Summing over all $d \in \{0, \dots, \log_2 n_0\}$ yields that the total stability of ALG from $t = 0$ until $t = n_0$ is lower bounded by $n_0(\log_2 n_0 + 1)$. ◀

Lemma 15 concludes our proof of Theorem 10 since we have shown that any algorithm that is an $\alpha(n)$-approximation for some function α cannot be $o(\log n)$-stable. We remark that our proof is even stronger than what we state in Theorem 10 since Lemma 15 lower bounds the *amortized* stability of any approximation algorithm.

▶ Remark. The lower bound from Theorem 10 also holds in the bipartite setting. Consider the exact same construction. Let $\mu(k)$ be the number of $i \in \mathbb{N}$ such that $k(i) = 1$, where $k \in \mathbb{N}$. Then $p(k)$ is a server if $\mu(k)$ is even, otherwise it is a client. Note that $\mu(k+1) = \mu(k) + 1$ if k is even, so exactly one server and one client arrives in each event. Moreover, Lemma 12 still holds using the same matching M from the proof. For any edge $p(k)p(\ell) \in M$ with $k < \ell$, we have $\mu(\ell) = \mu(k) + 1$, hence $p(k)$ is a server and $p(\ell)$ is a client or vice-versa, thus M is a perfect bipartite matching. The rest of the proof follows identically.

4 Concluding remarks

We studied the stability of online algorithms for the minimum-weight perfect matching problem in $\mathbb{R}^1$. We presented an $O(\sqrt{n})$-stable algorithm with approximation ratio 2, which we showed to be optimal among any algorithm with sublinear stability. We conjecture that $\Omega(\sqrt{n})$ is necessary to achieve approximation ratio 2. We also proved that any algorithm with $o(\log n)$ stability must have an unbounded approximation ratio. It would be interesting to see what approximation ratio can be achieved with an $f(n)$-stable algorithm for $f(n)$ in the range from $\Theta(\log n)$ to $\Theta(\sqrt{n})$. We conjecture that it is possible, for any integer constant $\alpha > 1$, to deterministically maintain a $(2\alpha - 1)$-approximation with an $O(n^{1/\alpha})$-stable algorithm similar to what Goranci et al. [6] achieve using randomization in the bipartite setting.

References

1 Mark de Berg, Arpan Sadhukhan, and Frits Spieksma. Stable approximation algorithms for the dynamic broadcast range-assignment problem. *SIAM Journal on Discrete Mathematics*, 38(1):790–827, 2024. doi:10.1137/23M1545975.

2 Mark de Berg, Arpan Sadhukhan, and Frits Spieksma. Stable approximation algorithms for dominating set and independent set. *SIAM Journal on Discrete Mathematics*, 39(2):921–945, 2025. doi:10.1137/24M1644079.

3 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. doi:10.6028/JRES.069B.013.

4 Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965. doi:10.4153/CJM-1965-045-4.

5 Harold N. Gabow. Data structures for weighted matching and nearest common ancestors with linking. In *Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 434–443. SIAM, 1990. URL: http://dl.acm.org/citation.cfm?id=320176.320229.

6 Gramoz Goranci, Peter Kiss, Neel Patel, Martin P. Seybold, Eva Szilagyi, and Da Wei Zheng. Fully dynamic euclidean bi-chromatic matching in sublinear update time. In *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*, volume 267 of *Proceedings of Machine Learning Research*. PMLR / OpenReview.net, 2025. URL: https://proceedings.mlr.press/v267/goranci25a.html.

7 Varun Gupta, Ravishankar Krishnaswamy, and Sai Sandeep. Permutation strikes back: The power of recourse in online metric matching. In *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques*, volume 176 of *LIPIcs*, pages 40:1–40:20, 2020. doi:10.4230/LIPIcs.APPROX/RANDOM.2020.40.

8 Bala Kalyanasundaram and Kirk Pruhs. Online weighted matching. *J. Algorithms*, 14(3):478–488, 1993. doi:10.1006/JAGM.1993.1026.

9 Samir Khuller, Stephen G. Mitchell, and Vijay V. Vazirani. On-line algorithms for weighted bipartite matching and stable marriages. *Theor. Comput. Sci.*, 127(2):255–267, 1994. doi:10.1016/0304-3975(94)90042-6.

10 Jannik Matuschke, Ulrike Schmidt-Kraepelin, and José Verschae. Maintaining perfect matchings at low cost. In *46th International Colloquium on Automata, Languages, and Programming*, volume 132 of *LIPIcs*, pages 82:1–82:14, 2019. doi:10.4230/LIPIcs.ICALP.2019.82.

11 Nicole Megow and Lukas Nölke. Online metric matching on the line with recourse. *Algorithmica*, 87(6):813–841, 2025. doi:10.1007/S00453-025-01299-8.

12 Krati Nayyar and Sharath Raghvendra. An input sensitive online algorithm for the metric bipartite matching problem. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 505–515. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.53.

13 Enoch Peserico and Michele Scquizzato. Matching on the line admits no o($\sqrt{\log n}$)-competitive algorithm. *ACM Trans. Algorithms*, 19(3):28:1–28:4, 2023. doi:10.1145/3594873.

14 Sharath Raghvendra. A robust and optimal online algorithm for minimum metric bipartite matching. In *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques*, volume 60 of *LIPIcs*, pages 18:1–18:16, 2016. doi:10.4230/LIPIcs.APPROX-RANDOM.2016.18.

15 Sharath Raghvendra. Optimal analysis of an online algorithm for the bipartite matching problem on a line. In *34th International Symposium on Computational Geometry, SoCG*, volume 99 of *LIPIcs*, pages 67:1–67:14, 2018. doi:10.4230/LIPIcs.SOCG.2018.67.

16 Pravin M. Vaidya. Geometry helps in matching. *SIAM J. Comput.*, 18(6):1201–1225, 1989. doi:10.1137/0218080.

17 Kasturi R. Varadarajan. A divide-and-conquer algorithm for min-cost perfect matching in the plane. In *Proc. 39th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 320–331. IEEE Computer Society, 1998. doi:10.1109/SFCS.1998.743466.