Online maximum matching with recourse

Spyros Angelopoulos[1] · Christoph Dürr[1] · Shendan Jin[1]

Abstract

We study the online maximum matching problem in a model in which the edges are associated with a known recourse parameter k. An online algorithm for this problem has to maintain a valid matching while edges of the underlying graph are presented one after the other. At any moment the algorithm can decide to include an edge into the matching or to exclude it, under the restriction that at most k such actions per edge take place, where k is typically a small constant. This problem was introduced and studied in the context of general online packing problems with recourse by Avitabile et al. (Inf Process Lett 113(3):81–86, 2013), whereas the special case $k = 2$ was studied by Boyar et al. (Proceedings of the 15th workshop on algorithms and data structures (WADS), pp 217–228, 2017). In the first part of this paper we consider the *edge arrival* model, in which an arriving edge never disappears from the graph. Here, we first show an improved analysis on the performance of the algorithm AMP of Avitabile et al., by exploiting the structure of the matching problem. In addition, we show that the greedy algorithm has competitive ratio $3/2$ for every even k and ratio 2 for every odd k. Moreover, we present and analyze an improvement of the greedy algorithm which we call L-GREEDY, and we show that for small values of k it outperforms the algorithm AMP. In terms of lower bounds, we show that no deterministic algorithm better than $1 + 1/(k − 1)$ exists, improving upon the known lower bound of $1 + 1/k$. The second part of the paper is devoted to the *edge arrival/departure model*, which is the fully dynamic variant of online matching with recourse. The analysis of L-GREEDY and AMP carry through in this model; moreover we show a lower bound of $(k^2 − 3k + 6)/(k^2 − 4k + 7)$ for all even $k \geq 4$. For $k \in \{2, 3\}$, the competitive ratio is $3/2$.

Keywords Matching · Online algorithms · Competitive analysis · Recourse

✉ Spyros Angelopoulos
  spyros.angelopoulos@lip6.fr

Extended author information available on the last page of the article

1 Introduction

In the standard framework of online computation, the input to the algorithm is revealed incrementally, i.e., as a sequence of *requests*. For each such requested input item, the online algorithm must make a decision that is typically *irrevocable*, in the sense that the algorithm commits, in a permanent manner, to the decision associated with the request. More precisely, the algorithm may not alter any previously made decisions while considering later requests. This rather stringent constraint is meant to capture what informally can be described as "the past cannot be undone". Equally significantly, it is at the heart of adversarial arguments that can be used to argue that its performance, measured by means of the *competitive analysis* framework (Borodin and El-Yaniv 1998), cannot be improved beyond a certain bound.

Nevertheless, there are real-life applications in which some (limited) rearrangement of the online solution during the execution of the algorithm may be doable, or even requisite. For instance, online call admission protocols may sporadically reconfigure the virtual paths assigned in the network. For a different example, in online scheduling (or resource allocation) problems, it may be permissible for a job to be transfered to a processor other than the one specified by the original decision associated with the job. Clearly, a trade-off is to be found between the guaranteed competitive ratio and the cost of re-optimizing the current solution. Different approaches to this objective have been considered. One such approach has studied the minimum total re-optimization cost required in order to maintain an optimal solution, see Bernstein et al. (2018). Another approach has focused on the best achievable competitive ratio when there is some bound on the allowed re-optimization, such as in Avitabile et al. (2013), and it is the main model we consider in this work.

More specifically, we study the *online maximum matching* problem, in which the objective is to maintain a vertex-disjoint edge set of maximum cardinality in a given graph. Here, the request sequence consists of the edges of the graph, and for each revealed edge, the online algorithm must decide whether to include it in the matching or not. In the standard model, each request has to be immediately either accepted in the matching, or rejected, and this decision cannot be revoked in the future. We consider a generalization of this model, in which the algorithm can switch between accepting and rejecting an edge that has already appeared, but is allowed up to k such modifications per edge, where k is the *recourse parameter* that is known to the algorithm.[1] The following section provides the formal definition of the problem, as well as a motivating application.

1.1 Problem definition and motivation

We define the *online matching problem with k edge-recourse*, for given $k \in \mathbb{N}^*$. The request sequence is a permutation of the edge set E in a graph $G = (V, E)$. The online algorithm knows the set V, as well as the parameter k, but not the set E. At each point in time, the online algorithm must maintain a matching $M \subseteq E'$, where

---

[1] For consistency with other recourse models, we define the initial default state of an edge as *rejected*, and therefore rejecting a newly arriving edge does not count as decision modification. See also Sect. 1.3.1.

$E' \subseteq E$ is the set of currently revealed edges. Specifically, for each revealed edge e in the sequence, the online algorithm either *accepts* e, by adding it in its matching, or *rejects* it. In addition, the algorithm must obey the edge recourse constraint, which is defined as follows. Every edge has an integer type, which is set to 0 upon the arrival of an edge. Whenever the algorithm decides either to include an edge e to its matching or to remove it, the type of e is increased. The algorithm can perform these operations at any point in time subject to the constraint that no edge type exceeds k.

The objective of the problem is to design an online algorithm of minimum *competitive ratio* which is defined as the worst–case ratio, over all request sequences σ, of the cardinality of an optimal offline matching to the cardinale of the matching produced by the online algorithm. See Sect. 1.2 for a more elaborate discussion on this measure.

We make a distinction between two settings concerning the revealed edges. In the first setting, which we call the *edge arrival* setting, once an edge appears (as a request) it is guaranteed to be part of the input graph G. In other words edges may only arrive, but never depart. A more general setting is the one in which edges may not only arrive (in the form of a request), but may also disappear adversarially (subsequently to their appearance). More specifically, at each request, the online algorithm is also informed about the subset of edges that are no longer part of the input graph since the last request. We call this setting the *edge arrival/departure* setting. Note that in this model, a given edge e may appear and disappear several times, and the number of such events is unrelated to the recourse parameter k. This setting is motivated by similar models that have been studied in the context of the online Steiner tree problem (Gupta and Kumar 2014), and models the fully dynamic variant of the problem at hand.

Last, in what concerns the edge arrival/departure setting, we make a further distinction concerning the edge departures. In the *full departure* model, the adversary is allowed to delete any edge in the graph, and thus also any edge that may have been provisionally accepted by the online algorithm. We show that this model is quite restrictive, since it yields excessive power to the adversary. We thus also study the *limited departure* model, in which the adversary may delete only edges not currently accepted by the online algorithm.

To motivate the problem and the various models, consider the following application related to resource allocation. Suppose that we have a set of tasks T and a set of workers W, as well as a bipartite graph $G = (W, T, E)$. An edge $(w, t) \in E$ between a worker w and a task t signifies that w is qualified to work on t, and we assume that a task can be assigned only to one worker and, vice versa, a worker can only be occupied in one task. A maximum matching in G describes a maximum assignment of tasks to workers. In the standard online version of the problem (with no recourse), the input consists of pairs of the form (w, t), and once an algorithm assigns worker w to task t (namely, accepts the edge (w, t)) it cannot assign w to a task other than t or t to a worker other than w. In contrast, the online matching problem with k edge-recourse captures the application in which w may be assigned in an "on/off" manner to task t up to k times. Here k bounds the willingness of workers to be reassigned to tasks.

Last, the edge arrival/departure model can capture the dynamic situation in which the compatibility of tasks and worker skills can change over time. Under the full departure model, there are no constraints on such changes. In contrast, the limited departure model stipulates that if worker w remains assigned to task t by the online

algorithm, then w does not lose his qualification for t, in the sense that the worker has a continual occupation with the said task and maintains the required skills for the task. However, once the online algorithm decides to remove worker w from task t (i.e., the online algorithm provisionally rejects edge (w, t)), then the worker might lose its qualification for the task over time.

1.2 Competitive analysis

Given an online algorithm ALG, and a request sequence σ, we use the standard notation $\text{ALG}(\sigma)$ to denote the output of ALG on sequence σ. In the context of the matching problem, we will denote by $|\text{ALG}(\sigma)|$ the *value* of the matching output by ALG on sequence σ, namely the cardinality of its matching. We will denote by OPT the offline optimal algorithm that has knowledge of the request sequence, and hence $|\text{OPT}(\sigma)|$ denotes the cardinality of the offline optimal matching.

An online algorithm ALG for the maximum cardinality matching problem is said to be *asymptotically c-competitive* if there is a constant d such that $|\text{ALG}(\sigma)| \geq |\text{OPT}(\sigma)|/c - d$, for all request sequences σ. If $d = 0$ the algorithm is called *strictly c-competitive*. Note that some previous work on the matching problem has used the reciprocal ratio. The smallest c for which an online algorithm ALG is c-competitive is called the *asymptotic competitive ratio* of ALG. The strict competitive ratio is defined similarly. If it so happens that this minimum value does not exist, the competitive ratio is actually defined by the corresponding infimum. In this setting an *upper bound* on the (strict or not) competitive ratio establishes the performance guarantee of an online algorithm, whereas a *lower bound* is a negative result.

Both upper and lower bounds in this work are shown for the strict competitive ratio. This implies that the upper bounds carry over to the more general definition, but this generalization does not necessarily hold for the lower bounds. We emphasize, however, that the known lower bounds for edge-bounded recourse problems in Avitabile et al. (2013) and Boyar et al. (2017) are likewise expressed in terms of the strict competitive ratio. This is due, perhaps, to difficulties in applying techniques that extend the lower bounds to the standard definition of the competitive ratio that are inherent to the recourse setting, and which do not arise in the traditional online framework of irrevocable decisions. Specifically, it is not obvious how to use techniques based on multiple copies of an adversarial instance in order to lower-bound the performance of any online alghorithm, although this may be possible for specific online algorithms (see, e.g., Lemma 2). For convenience, we will henceforth refer to the strict competitive ratio as simply the "competitive ratio".

For convenience of notation, we will omit the request sequence σ when it is implied from context, or when it is not relevant. Thus, with a slight, but standard abuse of notation, we will denote by ALG both the algorithm and its output.

1.3 Related work

Several online combinatorial optimization problems have been studied under recourse settings. The broad objective is to quantify the trade-off between the competitive ratio

and a measure on the modifications allowed on the solution. Some representative examples include online problems such as minimum spanning trees and TSP (Megow et al. 2016), Steiner trees (Gupta and Kumar 2014; Gu et al. 2016), knapsack problems (Han and Makino 2009; Iwama and Taketomi 2002), assignment problems in bipartite unweighted graphs (Gupta et al. 2014), and general packing problems (Avitabile et al. 2013). In the remainder of this section we review work related to online maximum matching.

1.3.1 Online matching, with and without recourse: models

Concerning online matching, two different request models have been studied in the past. In the *vertex arrival model*, vertices arrive in online fashion, revealing, at the same time, the edges incident with previously arrived vertices. This model has mainly been considered for bipartite graphs, with left side vertices arriving online, and right side vertices being initially known [(see the survey Mehta (2013)]. In the *edge arrival model* the edges arrive online in arbitrary order, revealing at the same time incident vertices.[2]

In the standard online model, every request (either vertex, or edge, depending on the request model) is served in an irrevocable manner. In contrast, for online matching with recourse and edge arrivals, several models have been proposed that relax the irrevocable nature of a decision. In the *late reject* model (Boyar et al. 2017), which is also called the *preemptive model* (Chiplunkar et al. 2015), an edge can be accepted only upon its arrival, but can be later rejected. The problem we study in this work, namely online matching with k edge-recourse was introduced in Avitabile et al. (2013). Boyar et al. (2017) refer to this model for $k = 1$ as the *late accept* model, and for $k = 2$ as the *late accept/reject* model. Clearly, the competitive ratio is monotone in k. Figure 1 provides an illustration of the algorithm's actions under the different models.

1.3.2 Online matching, with and without recourse: known results

For online maximum matching without recourse, and for the vertex arrival model, the seminal work of Karp et al. (1990) gave a randomized online algorithm with competitive ratio $e/(e - 1)$ in the vertex arrival model together with a matching lower bound on any online algorithm. In contrast, for the edge arrival model and the randomized competitive ratio, Buchbinder et al. (2017) showed a lower bound of $(3 + 1/\varphi^2)/2$ as well as an upper bound of 1.8 for the special case of forests, where φ is the golden ratio.

It is well known that any inclusion-wise maximal matching has cardinality at least half of the optimal maximum cardinality matching. From this it follows that the greedy online algorithm, which accepts an edge as long as it can be added to the current matching, has competitive ratio at most 2, which in the standard model is optimal among all deterministic online algorithms.

---

[2] We emphasize that in our work we consider the maximum cardinality matching problem; some previous work (with or without recourse) has considered the generalized *weighted matching* problem, in which each edge has a weight and the objective is to maximize the weight of edges in the matching.
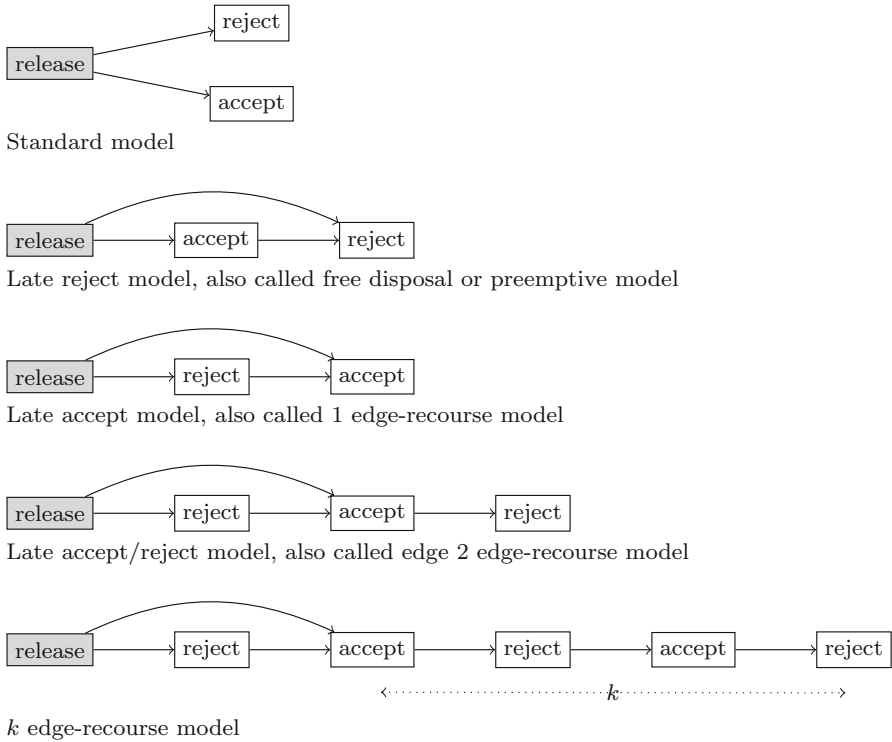
Fig. 1 Illustration of the actions of an online matching algorithm under the different edge-arrival models with recourse

Late reject In the vertex arrival model, the greedy algorithm achieves trivially the competitive ratio of 2, which is optimal for all deterministic online algorithms. The situation differs in the edge arrival model. Epstein et al. (2013) showed that for online weighted matching, the deterministic competitive ratio is exactly $3 + 2\sqrt{2} \approx 5.828$, as the upper bound of McGregor (2005) matches the lower bound of Varadaraja (2011). The same paper (Epstein et al. 2013) shows that the randomized competitive ratio is between $1 + \ln 2 \approx 1.693$ and 5.356. Chiplunkar et al. (2015) presented a randomized 28/15-competitive algorithm based on a primal-dual analysis.

k edge-recourse This model was introduced and studied by Avitabile et al. (2013) for the edge arrival setting, in the context of a much broader class of online packing problems. They gave an algorithm, which we call AMP, that combines doubling techniques with optimal solutions to offline instances of the problem, and which has competitive ratio $1 + O\left(\frac{\log k}{k}\right)$ (see Sect. 2.1 for an analysis of AMP). We note that this result is formulated in Avitabile et al. (2013) in a "dual" setting. More precisely, Avitabile et al. (2013) asks the question: how big should the edge budget k be such that there is a $(1 + \varepsilon)$-competitive online algorithm that makes at most k changes per edge? They showed that $k = O(\ln(1/\varepsilon)/\varepsilon)$ suffices. On the negative side, they showed that no

randomized algorithm can be better than $1 + 1/(9k - 1)$-competitive; we note also that their construction implies a lower bound of $1 + 1/k$ for all deterministic algorithms.

Boyar et al. (2017) showed that the deterministic competitive ratio is 2, if $k = 1$, and $3/2$, if $k = 2$, and this optimal competitive ratio is achieved by the greedy algorithm. Moreover Boyar et al. (2017) studied several other problems for a value of the recourse parameter equal to 2, such as independent set, vertex cover and minimum spanning forest.

Minimizing recourse   Bernstein et al. (2018) studied a different recourse model in which the algorithm has to maintain an optimal matching. More specifically, recourse here is expressed in terms of the total number of times edges enter or leave the algorithm's matching. They considered the setting of a bipartite graph and the vertex arrival model and showed that a simple greedy algorithm achieves optimality using $O(n \log^2 n)$ replacements, where n is the number of nodes in the arriving bipartition, whereas the corresponding lower bound for any replacement strategy is $\Omega(n \log n)$.

1.4 Contribution of this work

In the first part of this work, we study the online matching problem with edge k-bounded recourse under the edge arrival model. For this problem, we provide improvements on both upper and lower bounds on the competitive ratio. First, we revisit the doubling algorithm of Avitabile et al. (2013) that was originally analyzed in the general context of online packing problems. We give a better analysis, specifically for the problem at hand, that uses concepts and ideas related to the matching problem; we also show that the AMP algorithm has competitive ratio $1 + O(\frac{\log k}{k})$. On the negative side, we show that no deterministic algorithm is better than $(1 + 1/(k - 1))$-competitive, improving upon the known bound of $1 + 1/k$ of Avitabile et al. (2013).

At first sight these improvements may seem marginal; however one should take into consideration that k is typically a small parameter, and thus the improvements are by no means negligible. In this spirit, we propose and analyze a variant of the greedy algorithm which we call L-Greedy. This algorithm applies, at any step, augmenting paths as long as their length is at most $2L + 1$. We show that for a suitable choice of L, this algorithm is $1 + O(1/\sqrt{k})$-competitive. While this algorithm is thus not superior to AMP for large k (and more specifically, to its improved analysis in the context of the matching problem), for small k (and in particular, for $k \leq 20$) it does achieve an improved competitive ratio, see Fig. 2. Moreover, we extend a result of Boyar et al. (2017) that showed that the greedy algorithm is $\frac{3}{2}$-competitive for $k = 2$ to all even k (for odd k, the competitive ratio is 2).

In terms of techniques, we analyze both AMP and L-Greedy using amortization arguments in which the profit of the algorithms is expressed in terms of weights appropriately distributed over nodes in the graph. We achieve these improvements by exploiting properties of augmenting paths in matching algorithms.

The second part of the paper is devoted to the *edge arrival/departure model*, which is the fully dynamic variant of the online matching problem. First, we observe that the analysis of L-Greedy and AMP carries through in this model as well. On the
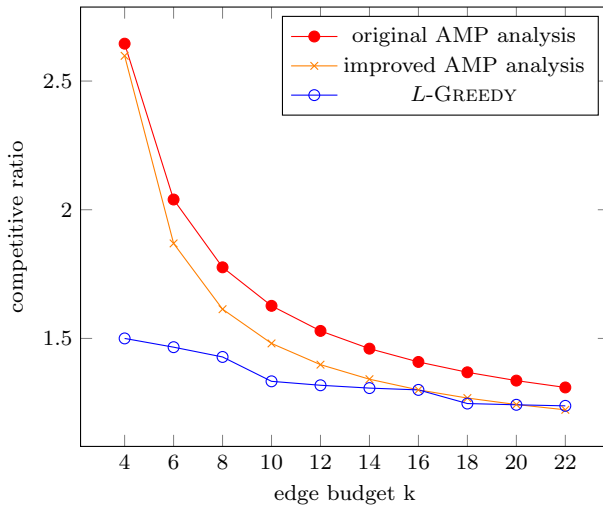
Fig. 2 Comparison of the competitive ratios of the algorithm AMP and the algorithm *L*-GREEDY

negative side, we show a lower bound of $(k^2 - 3k + 6)/(k^2 - 4k + 7)$ for all even $k \geq 4$. For $k \in \{2, 3\}$, the competitive ratio is $3/2$. We obtain these lower bounds by modeling the game between the algorithm and the adversary as a game played over *strings* of numbers 0 up to k. These strings represent alternating paths, and each number represents how many times the algorithm has modified its decision on the corresponding edge. This provides a simpler combinatorial aspect to the game played between the adversary and the algorithm.

We note that, for the analysis of AMP and of L-GREEDY, we assume that k is even. This assumption is borrowed from Avitabile et al. (2013) and is required for the analysis. Of course for odd $k \geq 3$ these algorithms can be run with budget $k - 1$, providing a valid upper bound on the competitive ratio. Note that our lower bound in the arrival model holds for all values of k.

1.5 Preliminaries

A *matching* in a graph $G = (V, E)$ is a set of edges $M \subseteq E$ with disjoint endpoints. A vertex $v \in V$ is said to be *matched* by M if there is an edge $e \in M$ incident to v, and is *unmatched* otherwise. A key concept in maximum matching algorithms is the notion of an *augmenting alternating path*, or simply *augmenting path*. A path P in G is a sequence of vertices $v_1, v_2, \ldots, v_\ell$ for some length $\ell \geq 2$, such that $(v_i, v_{i+1}) \in E$ for all $i = 1, \ldots, \ell - 1$. It is said to be *alternating with respect to M* if every other edge of P belongs to M. Moreover, an alternating path is *augmenting* if the first and the last vertex in the path is unmatched by M. Applying P to M consists in removing from M the edges in $M \cap P$ and adding the edges in $P \setminus M$. The resulting matching has cardinality $M + 1$, and every previously matched vertex remains matched.

We define some concepts that will be useful in the analysis of algorithms throughout the paper. We will associate each edge with a *type* which is an integer in $[0, k]$. An edge

is of type i if it has undergone i decision flips by the algorithm. Hence, for an edge of type k, where k is the recourse budget, its decision has been finalized, and cannot change further; we call such an edge *blocked*. The type of a path P is defined by the sequence of the types of its edges, and to make this concept unambiguous, we choose between the two orientations of the path the one that results in the lexicographically minimal such sequence. Note that when the algorithm applies some augmenting path P to its current matching M, then the type of every edge in P is increased by 1. Moreover, the two extreme edges of an augmenting path are of type 0, because the endpoints of P are unmatched. We will call a path *blocked* if it contains a blocked edge.

Given a request sequence σ, let $G(\sigma)$ denote the graph induced by σ. For an online algorithm ALG, and a connected component C of $G(\sigma)$, we define the *local ratio* in C as the ratio between the cardinality of the optimum matching and the cardinality of $ALG's$ matching, restricted to edges in C.

2 Online matching in the edge arrival model

2.1 The algorithm AMP

We study the performance of an algorithm proposed by Avitabile et al. (2013) for the more general *online set packing problem*. In this problem sets arrive online, and the objective is to maintain a collection of disjoint sets that has maximum cardinality. More specifically, Avitabile et al. (2013) proposed a *doubling algorithm* which is defined for even k only. The algorithm has a parameter $r > 1$ and there is a decision variable for every set which can be changed at most k times. The algorithm works in phases, sequentially numbered by an integer p. Initially $p = 0$, and the algorithm's current solution is $\text{AMP}_0 = \emptyset$. Let ℓ be the largest integer such that the optimal solution has value at least r^ℓ, where ℓ is defined to be $-\infty$ if the optimal solution is empty. Whenever this value increases, the algorithm starts a new phase. We define $\ell(p)$ as the value of ℓ during phase p, and thus

$$\ell(p) + i \leq \ell(p + i) \tag{1}$$

for every positive integer i. At the beginning of a new phase, all decision variables that have been changed fewer than k times are set as in OPT, resulting in the current solution AMP_p (note that the algorithm crucially depends on k being even in order to produce a feasible solution). In between phases, AMP does not make any changes to its current solution.

Avitabile *et al* show that the value $|AMP|$ of the solution returned by the algorithm is at least

$$\left(1 - \left(\frac{r-1}{r} + \frac{r}{r^k(r-1)}\right)\right)|OPT|,$$

which implies that

$$\frac{|OPT|}{|AMP|} \leq \frac{1}{1 - \left(\frac{r-1}{r} + \frac{r}{r^k(r-1)}\right)} = \frac{r^k(r-1)}{r^{k-1}(r-1) - r}.$$

Thus, for a given $r > 1$, the competitive ratio of AMP is at most

$$\frac{r^k(r-1)}{r^{k-1}(r-1) - r}.$$

Let r_0 denote the root of $r^{k-1}(r-1) - r = 0$. If $r < r_0$, then for all $k \geq 4$, we have $r^{k-1}(r-1) - r < 0$ since this function is increasing in r. We conclude that the competitive ratio of AMP is upper bounded by

$$\min_{r>r_0} \frac{r^k(r-1)}{r^{k-1}(r-1) - r}. \tag{2}$$

We will show how to obtain an improved analysis of this algorithm in the context of the matching problem. Since we know optimal algorithms for $k = 1, 2$ by Boyar et al. (2017) and, since we can show optimality of the greedy algorithm for $k = 3$ (see Sect. 2.4), we can assume, for the purposes of the analysis, that $k \geq 4$. We begin by a restatement of the update phase that will help us exploit the structure of solutions obtained via augmenting paths. More specifically, on every edge arrival, the algorithm updates a current optimal solution OPT. At the beginning of a new phase, the algorithm produces a matching AMP_p obtained from AMP_{p-1} as follows: every edge $e \in AMP_{p-1} \setminus OPT$ is removed from the current matching, and every edge $e \in OPT \setminus AMP_{p-1}$ which is of type strictly smaller than k is added to the current matching. Note that edges incident with e have been removed, hence AMP_p is indeed a matching. Also note that all edges added or removed by the algorithm have their type increased by one.

Since AMP_{p-1} and OPT are matchings, their symmetric difference, excluding type k edges, consists of alternating cycles and alternating paths which can be of even or odd length. This means that the algorithm simply applies at the beginning of every phase all those alternating paths and cycles.

Let OPT_p denote the matching produced by OPT as phase p is about to begin. From the statement of AMP, we obtain the following series of inequalities, for every phase p.

$$r^{\ell(p)} \leq |OPT_p| \leq |OPT| < r^{\ell(p)+1}. \tag{3}$$

For any given phase $p \geq 1$, we aim to bound the ratio $|OPT|/|AMP_p|$, since this will allow us to bound the competitive ratio of AMP, for a worst–case choice of p. Note that the type of an edge increases by at most 1 with each phase. Hence, in the beginning of the k first phases AMP "synchronizes" with OPT as there are no blocked edges yet, and as a result during these phases the ratio $|OPT|/|AMP_p|$ does not exceed r, from (3).

For the remaining phases we need the following argument.

Proposition 1 *For even k and any phase $p \geq k+1$, AMP maintains a matching AMP_p of cardinality at least $r^{\ell(p)} - r^{\ell(p-k+1)+1}$.*

Proof We denote by the *type of a vertex* v the maximum type of the edges incident with v and by $n_{i,p}$ the number of vertices of type i in phase p. With every phase change the type of a vertex increases at most by 1. Hence every vertex of type k in phase p had positive type in phase $p - k + 1$. Thus

$$n_{k,p} \leq \sum_{i=1}^{k} n_{i,p-k+1} \leq 2 \cdot |\text{OPT}_{p-k+1}|,$$

where the last inequality uses the fact that the left hand side counts the number of vertices matched by the algorithm. In phase p, the difference between the cardinality of the optimal matching and the cardinality of AMP's matching is at most the number of blocked augmenting paths, and each of them contains at least two type k vertices. Hence,

$$
\begin{aligned}
|AMP_p| &\geq |\text{OPT}_p| - \frac{1}{2} \cdot n_{k,p} \\
&\geq |\text{OPT}_p| - |\text{OPT}_{p-k+1}| \\
&> r^{\ell(p)} - r^{\ell(p-k+1)+1},
\end{aligned}
$$

where the last inequality follows from (3). $\qquad\square$

Combining Proposition 1 with the bounds (3) we obtain the following.

Proposition 2 *The competitive ratio of AMP for $k \geq 4$ is upper bounded by the expression*

$$\min_{r>1} \frac{r^k}{r^{k-1} - r}. \tag{4}$$

Proof Consider an arbitrary phase p and a fixed parameter $r > 1$. The expression $r^k/(r^{k-1} - r)$ is at least r. As observed earlier, at the end of the first k phases, the competitive ratio is at most r, hence the proof holds for $p \leq k$. For the remaining case $p \geq k + 1$, we have that

$$
\begin{aligned}
\frac{|\text{OPT}|}{|\text{AMP}_p|} &\leq \frac{r^{\ell(p)+1}}{r^{\ell(p)} - r^{\ell(p-k+1)+1}} &&\text{(From (1))} \\
&\leq \frac{r^{\ell(p)+1}}{r^{\ell(p)} - r^{\ell(p)-k+2}} \\
&= \frac{r}{1 - r^{-k+2}} \\
&= \frac{r^k}{r^{k-1} - r}.
\end{aligned}
$$

$\qquad\square$

Next, we show that Proposition 2 can yield an improved analysis of AMP over the original bound (2).

Proposition 3 *For all even $k \geq 4$, we have that the competitive ratio as expressed by (2) is at least the expression (4).*

Proof For $k = 4$, we obtain numerically that (2) is 2.64526 whereas (4) is 2.59808.

For even $k \geq 6$, first we show that the minimizer (for $r > 1$) of

$$\frac{r^k(r-1)}{r^{k-1}(r-1) - r}$$

is between r_0 and 2. The derivative of the above expression is

$$\frac{(r-1)^2 r^{2k} + (k-1-kr)r^{k+2}}{\left(r^2 - (r-1)r^k\right)^2},$$

which we claim to be positive for any $r \geq 2$. This follows by the inequality $r^{k-3} \geq k$ which holds for any $r \geq 2$ and $k \geq 6$.

Similarly, for even $k \geq 6$, we show that the minimizer (for $r > 1$) of

$$\frac{r^k}{r^{k-1} - r}$$

is between 1 and 2. The derivative in r of the above expression is

$$\frac{r^k \left(r^k - (k-1)r^2\right)}{\left(r^2 - r^k\right)^2},$$

which we claim to be positive for any $r \geq 2$. Again this follows by the inequality $r^{k-2} \geq k - 1$, which holds for any $r \geq 2$ and $k \geq 6$. The proof follows since for $1 < r \leq 2$, we have

$$\frac{r^k(r-1)}{r^{k-1}(r-1) - r} \geq \frac{r^k(r-1)}{r^{k-1}(r-1) - r(r-1)} = \frac{r^k}{r^{k-1} - r}.$$

$\square$

The following theorem concludes the asymptotic analysis of the performance of AMP.

Theorem 1 *For all even k, AMP has competitive ratio $1 + O(\frac{\log k}{k})$.*

Proof We first sketch a simple argument based on the Puiseux series expansion (Siegel 1988): this is a type of power series that allows fractional powers, as opposed to only integer ones (e.g., Taylor series). Let r denote the optimal choice of the parameter, namely the one that minimizes (4). By analyzing the derivative, it follows that $r =$

$(k-1)^{1/(k-2)}$, hence the competitive ratio is at most $\frac{(k-1)^{\frac{k-1}{k-2}}}{k-2}$, whose Puiseux series expansion at $k = \infty$ is $1 + \frac{\log k + 1}{k} + O(\frac{1}{k^2})$.

For completeness, we give a second proof that relies only on standard calculus. Let f be such that $r = 1 + f/k$, then since $r = (k-1)^{1/(k-2)}$, we have that the competitive ratio is at most

$$r \cdot \frac{k-1}{k-2} = \left(1 + \frac{f}{k}\right)\frac{k-1}{k-2} = 1 + O\left(\frac{f}{k} + \frac{1}{k}\right). \tag{5}$$

Suffices then to show that $f = O(\log k)$. Consider the function

$$g(x) = x^{k-2} - k + 1,$$

and note that r must be a root of g. We can rewrite $g(r)$ as

$$g(r) = \left(1 + \frac{f}{k}\right)^{k-2} - k + 1 = e^{(k-2)\ln(1+f/k)} - k + 1.$$

Hence, we have

$$g(r) \geq e^{(k-2)2f/(2k+f)} - k + 1.$$

Here we used the following logarithmic inequality (Love 1980) for $n \geq 0$:

$$\ln(1 + 1/n) \geq \frac{2}{2n+1}.$$

Suppose now that $f = 4 \ln k$. In this case, we claim that $(k-2)\frac{2f}{2k+f} \geq \frac{f}{2}$, or equivalently $2k \geq 8 + f$, which holds for sufficiently large k. Thus we have

$$g\left(1 + \frac{4\ln k}{k}\right) \geq e^{f/2} - k + 1 = k^2 - k + 1 > 0.$$

On the other hand, $g(0) = -k + 1 < 0$. As a result, since g is continuous, there exists $f \in [0, 4\ln k]$ such that $g(r) = 0$. Therefore, $f = O(\log k)$, which concludes the proof. $\qquad\square$

2.2 The algorithm GREEDY

We consider the algorithm GREEDY, which repeatedly applies an arbitrary augmenting path whenever possible. More precisely, let E denote the set of edges that have been revealed to the algorithm, and let e denote a newly arriving edge. Then as long as there is a non-blocked augmenting path in the graph $(V, E \cup \{e\})$, where V is the vertex set, GREEDY will apply such a path until, eventually, no such path any longer exists. Note, in particular, that GREEDY does nothing if no such path exists upon the arrival of e.

This algorithm achieves an upper bound of $3/2$ for $k = 2$, as shown in Boyar et al. (2017). We show that the same guarantee holds for all even k. In what concerns the
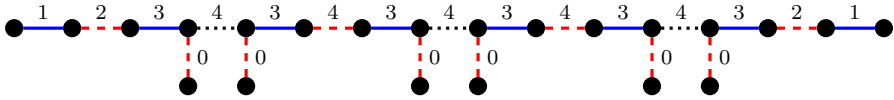
Fig. 3 Lower bound construction on the competitive ratio of GREEDY for $k = 4$ and $n = 2$. Edges are labeled by their type. Solid/blue edges depict the algorithm's matching, dashed/red edges depict the optimal matching, dotted/black edges belong to none of the matchings

lower bound, the idea is to force the algorithm to augment an arbitrarily long path in order to create a configuration with an arbitrarily large number of blocked augmenting paths of lengths 5, which have local ratio 3/2.

Proposition 4 *The competitive ratio of* GREEDY *is* 3/2 *if* k *is even, and* 2, *if* k *is odd.*

Proof First note that GREEDY has the property that every edge in the optimal matching has at least one endpoint matched by GREEDY. As a result, the competitive ratio is at most 2 for any k, and in particular for any odd k.

For even positive k we give a stronger upper bound of 3/2. Consider the symmetric difference of the matching produced by GREEDY and an optimal matching which consists of alternating cycles and paths. In each of these components the local ratio is 1 for the alternating cycles and alternating paths of even length. We claim that alternating paths of odd length have length ℓ at least 5, and therefore the local ratio is at most 3/2. To see this, observe that the case $\ell = 1$ corresponds to an edge with both endpoints unmatched, and GREEDY would have included it in its matching. In the case $\ell = 3$, the center edge has an odd budget, meaning that GREEDY would have applied this augmenting path.

The proof of the lower bound for even positive k consists of an instance on which GREEDY achieves the competitive ratio $3/2 - \epsilon$ for any small constant $\epsilon > 0$; see Fig. 3. Let n be a sufficiently large integer. First the adversary releases $2n + 1$ vertex disjoint edges, which the algorithm includes in its matching. Then these edges are connected with $2n + 2$ new edges to form an augmenting path of length $4n + 3$. From now on, each time the algorithm applies this augmenting path in its matching, the path is extended with one additional edge on each end, until there is at least one edge of type k on the path. At this point the edge types on the path form a sequence which starts with types from 1 to $k - 1$, then alternates between types k and $k - 1$ and finally ends with types from $k - 1$ to 1. To complete the instance, the adversary attaches a new edge on each endpoint of every second type k edge. As a result there are n augmenting paths of length 5, which are all blocked by a type k edge, together with 2 alternating paths of length k. The size of the matching produced by GREEDY is $2n + k$, whereas the optimal matching has size $3n + k$, showing a lower bound of 3/2.

For odd k the construction can be further strengthened. In the final step, the adversary attaches an edge on each endpoint of every type k edge, creating an arbitrary large number of blocked augmenting paths of length 3.

Note that the lower bounds hold even for the asymptotic competitive ratio, by repeating these constructions as in the proof of Lemma 2.   □

2.3 The algorithm *L*-GREEDY

How can the greedy algorithm be improved? As illustrated in the proof of Proposition 4, the greedy algorithm has inferior performance because it augments along arbitrarily long augmenting paths, therefore sometimes sacrificing edge budget for only a marginal increase in the matching size. A natural idea towards an improvement would be to apply only short augmenting paths, as they are more budget efficient. For technical reasons, we restrict the choice of augmenting paths even further.

We define the algorithm *L*-GREEDY for some given parameter L, which applies any non-blocked augmenting path of length at most $2L + 1$ that is in the symmetric difference between the current matching and some particular optimal matching OPT. The latter is updated after each edge arrival by applying an augmenting path to OPT.

To make the above more precise, let E denote the set of edges that have been revealed to the algorithm, and let e denote a newly arriving edge. First, we explain how the optimal matching is updated: If $\text{OPT}(E)$ denotes the optimal matching after all edges in E have been revealed, then $\text{OPT}(E \cup \{e\})$ is obtained by applying an augmenting path to $\text{OPT}(E)$. Then, *L*-GREEDY serves request e by consecutively applying non-blocked augmenting path of length at most $2L + 1$ that is in the symmetric difference between its current matching (prior to the arrival of e) and $\text{OPT}(E)$.

Note that *L*-GREEDY may not change its solution even if there is a short augmenting path in the current graph if it contains edges which are not in this particular optimal matching OPT. We will later optimize the parameter L as a function of k.

2.3.1 Analysis of *L*-GREEDY

We begin by observing that for $L = 0$ the algorithm collects greedily vertex disjoint edges without any recourse, which is precisely the behavior of GREEDY for $k = 1$ and has competitive ratio 2. For $L = 1$ the algorithm *L*-GREEDY applies only augmenting paths of length at most 3. In this case, the same argument as in the proof of Proposition 4 shows that the competitive ratio of *L*-GREEDY is 3/2.

In what follows we analyze the general case $L \geq 2$. To this end, we assign weights to vertices in a way that the total vertex weight equals the size of the current matching. Therefore, whenever the size of the matching is increased by 1, a total weight of 1 is distributed on the vertices along the augmenting path as follows. First, vertices in this path that were already matched receive a weight α, where $\alpha \geq 0$ is some constant that we specify later. Second, the two vertices on the endpoints of the augmenting path receive the remaining weight, that is $1/2 - \ell\alpha$, where $2\ell + 1$ is the length of the path. It follows, from this weight assignment, that every unmatched vertex has weight 0, that every matched vertex has weight at least $1/2 - L\alpha$, and that every endpoint of a type k edge has weight at least $1/2 - L\alpha + (k - 1)\alpha$.

Suppose that *L*-GREEDY reaches a configuration in which it cannot apply any augmenting path, as specified in its statement. We consider the symmetric difference between the matching produced by the algorithm and the optimal matching maintained internally by the algorithm. This symmetric difference consists of alternating paths and/or alternating cycles, and we will upper bound for each such component its local ratio. In particular, a component in the symmetric difference falls in one of the

following cases: Either it is an augmenting path of length $2\ell + 1 \leq 2L + 1$, or an augmenting path of length $2\ell + 1 > 2L + 1$, or an alternating cycle or alternating path of even length.

Case 1: Augmenting path of length $2\ell + 1 \leq 2L + 1$. Note that such a path contains at least one edge of type k, otherwise the algorithm would augment it. It follows that $\ell \geq 2$, since an augmenting path of length 1 is a single type 0 edge, and an augmenting path of length 3 has edge types respectively $0, t, 0$ for some odd t (and k is assumed to be even). The path contains 2ℓ matched vertices, and at least 2 of them are incident with a type k edge. Hence the total vertex weight is at least $2\ell \left(\frac{1}{2} - L\alpha\right) + 2(k - 1)\alpha$, and the local ratio of this component is at most

$$\frac{\ell + 1}{\ell - 2\ell L\alpha + 2(k - 1)\alpha}. \tag{6}$$

Case 2: Augmenting path of length $2\ell + 1 > 2L + 1$. Such a path contains 2ℓ matched vertices and therefore the local ratio is at most

$$\frac{\ell + 1}{\ell - 2\ell L\alpha}. \tag{7}$$

Case 3: Alternating cycle or path of even length. Such a component contains 2ℓ matched vertices and therefore the local ratio is at most

$$\frac{\ell}{\ell - 2\ell L\alpha},$$

which is dominated by (7). We obtain the following performance guarantee.

Theorem 2 *The competitive ratio of L-GREEDY with $L = \lfloor \sqrt{k - 1} \rfloor$ is at most*

$$\frac{k(L + 2) - 2}{(L + 1)(k - 1)} = 1 + O\left(\frac{1}{\sqrt{k}}\right),$$

for even $k \geq 6$ and at most $3/2$ for $k = 4$.

Proof We choose α so as to minimize the maximum of the local ratios, as defined by (6) and (7). Then for this choice of α we optimize L as stated in the theorem. Note that for $k = 4$, this leads to the choice $L = 1$, which we analyzed in the beginning of the section.

For $k \geq 6$, we have to minimize over α and L the maximum over ℓ of the two ratios given by (6) and (7). First we upper bound (7) as

$$\frac{\ell + 1}{\ell - 2\ell L\alpha} \leq \frac{L + 2}{L + 1 - 2L^2\alpha - 2L\alpha}, \tag{8}$$

where the inequality follows from $\ell \geq L + 1$ and the fact that the left hand side is decreasing in ℓ as can be seen by dividing both the numerator and denominator by ℓ.

In order to upper bound (6), we find its derivative in ℓ which is equal to

$$\frac{2\alpha(L + k - 1) - 1}{(\ell + 2\alpha(k - L\ell - 1))^2}.$$

This means that (6) is increasing or decreasing in ℓ depending on the sign of $2\alpha(L + k - 1) - 1$. Hence we distinguish two cases.

Case 1: $\alpha < 1/(2(L + k - 1))$

In this case (6) is decreasing in ℓ, and by $\ell \geq 2$ is at most

$$\frac{3}{2\alpha(k - 2L - 1) + 2}. \tag{9}$$

Subcase 1a: $k - 2L - 1 < 0$

In this case (9) and (8) are increasing in α and hence minimized at $\alpha = 0$. For this choice of α (9) is $3/2$, while (8) is $(L + 2)/(L + 1)$ which by $L \geq 2$ is less than $3/2$. In conclusion, the competitive ratio in Subcase 1a is at most $3/2$.

Subcase 1b: $k - 2L - 1 \geq 0$

In this case (9) is non-increasing in α while (8) is increasing in α. Hence the maximum of (9) and (8) is minimized at the equality of the expressions, which happens for

$$\alpha = \frac{L - 1}{2L^2 + 4k + 2kL - 4L - 4}.$$

It can readily be verified that this choice of α indeed satisfies the assumption of Case 1 for all even $k \geq 2$ and $L \geq 2$. The corresponding ratio is

$$\frac{L^2 + 2k + kL - 2L - 2}{(k - 1)(L + 1)}. \tag{R1b}$$

Case 2: $\alpha \geq 1/(2(L + k - 1))$ In this case (6) is increasing in ℓ, and at $\ell = L$ becomes

$$\frac{L + 1}{L + 2\alpha(k - L^2 - 1)}. \tag{10}$$

Subcase 2a: $k - L^2 - 1 < 0$ In this case, (10) and (8) are increasing in α. At $\alpha = 1/(2(L + k - 1))$, the former becomes

$$\frac{L + 1}{L + \frac{k - L^2 - 1}{L + k - 1}} = \frac{(L + 1)(L + k - 1)}{L(L + k - 1) + k - L^2 - 1}$$

$$= \frac{(L + 1)(L + k - 1)}{(L + 1)(k - 1)}$$

$$= \frac{L + k - 1}{k - 1}. \tag{11}$$

Similarly, (8) becomes

$$\frac{L+2}{L+1-\frac{L(L+1)}{L+k-1}} = \frac{(L+2)(L+k-1)}{(L+1)(L+k-1)-L(L+1)} \tag{R2a}$$

$$= \frac{(L+2)(L+k-1)}{(L+1)(k-1)},$$

which dominates (11) and therefore upper bounds the competitive ratio in Subcase 2a.

Subcase 2b: $k - L^2 - 1 \geq 0$ In this case, (10) is non-increasing in α while (8) is increasing in α. We have equality of the expressions for

$$\alpha = \frac{1}{2k(L+2)-4},$$

which implies that the competitive ratio in Subcase 2b is at most

$$\frac{k(L+2)-2}{(L+1)(k-1)}. \tag{R2b}$$

This concludes the case analysis. Thus it remains to optimize L.

For $k = 4$, the assumptions of the Subcases 1b and 2b are not valid. Hence, we have to choose L so to minimize the minimum of 3/2 and (R2a), which is larger than 3/2. This means that any choice $L \geq 2$ leads to a competitive ratio at most 3/2. Note that the same guarantee is obtained for $L = 1$.

Finally we consider $k \geq 6$ and choose L so to minimize the minimum of 3/2 and the ratios (R1b), (R2a), (R2b). We claim that (R1b) is dominated by 3/2. Indeed its derivative in k is

$$-\frac{L(L-1)}{(k-1)^2(L+1)} < 0.$$

Hence (R1b) is maximized for k, L such that $k - 2L - 1 = 0$, and this maximum value is precisely 3/2.

By comparing the numerators of the ratios (R1b), (R2a) and (R2b) the minimum ratio is (R2b). Its derivative in L is

$$-\frac{k-2}{(L+1)^2(k-1)} \leq 0.$$

Hence the minimum of (R2b) is attained at the upper bound for L given by the Subcase 2b assumption, namely

$$L = \lfloor \sqrt{k-1} \rfloor.$$

It follows that the competitive ratio is at most

$$\frac{k(\lfloor\sqrt{k-1}\rfloor+2)-2}{(\lfloor\sqrt{k-1}\rfloor+1)(k-1)} \leq 1 + \frac{2k+\sqrt{k-1}-1}{(\sqrt{k-1}+1)(k-1)} = 1 + O\left(\frac{1}{\sqrt{k}}\right).$$

□

We complete the analysis of L-GREEDY by showing that the upper bound is essentially tight.

Lemma 1 *For even $k \geq 4$, the strict competitive ratio of L-GREEDY with $L = \lfloor\sqrt{k-1}\rfloor$ is at least*

$$1 + \Omega\left(\frac{1}{\sqrt{k}}\right).$$

Proof For the proof we show that for any positive parameter $L \geq 3$ the (strict) competitive ratio of L-GREEDY is at least

$$\frac{3\lfloor\frac{L-1}{2}\rfloor+k-2}{L+k-3}. \tag{12}$$

The statement follows then by lower bounding this expression for $L = \lfloor\sqrt{k-1}\rfloor$.

We will show how to create an instance in which L-GREEDY has the competitive ratio expressed by (12). The adversarial construction is depicted in Fig. 4, and consists of the following steps. First the adversary releases $L-2$ vertex disjoint edges, denoted by $e_1, \ldots, e_{L-2}$, which the algorithm includes in its matching. Then these edges are connected with $L-1$ new edges (which are denoted by $f_1, f_2, \ldots, f_{L-1}$ in Fig. 4) so as to form an augmenting path P of length $2L-3$, which is short enough to ensure that the algorithm applies it. Note that the precise order in which the edges arrive is not important. At this stage the edges of P have alternating types 1 and 2, and by the choice of $L \geq 3$ there is at least one edge of type 2 in P.

In the next phase, the adversary releases edges from the sets A and B, in an interleaved manner; specifically, P is extended with two additional edges on each end, so as to form augmenting paths of length $2L-1$ and $2L+1$, until there are edges of type k on the path. At this point, the edge types on the medium part of length $2L-3$ form an alternating sequence of $k-1$ and k, and note that there are $k/2-1$ edges of type 2 and $k/2-1$ edges of type 1 in each set A and B. To complete the instance, the adversary attaches a new edge on each of the type $k-1$ edges in P, alternating between the left endpoint and on the right endpoint of P. As a result, all augmenting paths are blocked with a type k edge. The size of the matching produced by L-GREEDY is $L+k-3$, whereas the optimal matching has size at least $3\lfloor\frac{L-1}{2}\rfloor+k-2$, concluding the proof.

□

The previous lemma can be extended to the asymptotic competitive ratio, using a standard technique based on multiple copies of the adversarial instance.
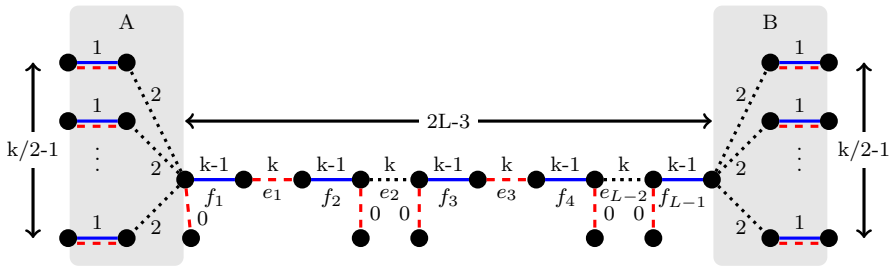
Fig. 4 Lower bound construction for even k and $L = 6$. Numbers on edges describe the edge types at the end of the algorithm's execution. Solid/blue edges depict the algorithm's matching, dashed/red edges depict the optimal matching and dotted/black edges belong to none of the matchings

Lemma 2 *For even k and $L = \lfloor\sqrt{k-1}\rfloor \geq 3$, the asymptotic competitive ratio of L-GREEDY is at least the expression* (12) *and hence is* $1 + \Omega(1/\sqrt{k})$.

Proof Assume that the asymptotic competitive ratio of L-GREEDY is R for R being strictly smaller than expression (12). Then by definition there exists a constant d such that

$$|L\text{-}\text{GREEDY}(\sigma)| \geq |\text{OPT}(\sigma)|/R - d$$

for all request sequences σ. Let σ' be the adversarial construction defined in the proof of Lemma 1. For an arbitrary positive integer p let σ be the result of repeating each edge of σ' p times, such that the resulting graph consists of p disjoint copies of the graph produced by σ'. By the above assumption we have

$$
\begin{aligned}
|L\text{-}\text{GREEDY}(\sigma)| &\geq |\text{OPT}(\sigma)|/R - d & \equiv \\
p \cdot |L\text{-}\text{GREEDY}(\sigma')| &\geq p \cdot |\text{OPT}(\sigma')|/R - d & \equiv \\
|L\text{-}\text{GREEDY}(\sigma')| &\geq |\text{OPT}(\sigma')|/R - d/p.
\end{aligned}
$$

Since d/p can be arbitrarily close to 0, this would mean that L-GREEDY is R-competitive, a contradiction. $\square$

2.4 Lower bound on the competitive ratio of deterministic algorithms

Boyar et al. (2017) show that the deterministic competitive ratio of the problem is 2 for $k = 1$ and $3/2$ for $k = 2$. We complete this picture by showing a lower bound of $1 + \frac{1}{k-1}$ for all $k \geq 3$. Note that the lower bound is tight for $k = 3$, as the algorithm GREEDY, which works by assuming that k is only 2, has competitive ratio $3/2$.

Theorem 3 *The deterministic competitive ratio of the online matching problem with k edge-recourse is at least $1 + \frac{1}{k-1}$ for all $k \geq 3$.*

Proof We consider three cases, namely the cases $k = 3$, k is even and at least 4, and finally k is odd and at least 5. For each case we present an appropriate adversarial argument.
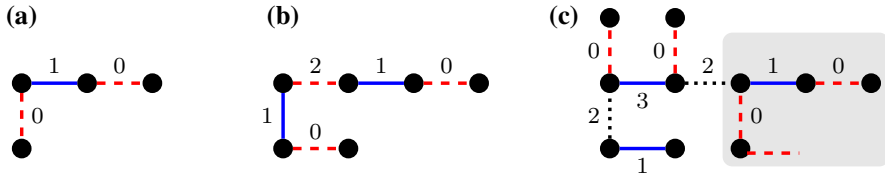
Fig. 5 Lower bound construction on the competitive ratio for the case $k = 3$

Case $k = 3$. Suppose, by way of contradiction, some algorithm claims a competitive ratio strictly smaller than $(3n + 2)/(2n + 2)$ for some arbitrary $n \geq 1$. The adversary releases a single edge, creating an augmenting path of length 1. Then the algorithm applies the augmenting path, which the adversary extends by appending one edge on each side, creating an augmenting path of type 0, 1, 0, as shown in Fig. 5a. Since the current competitive ratio is 2, the algorithm needs to apply this path, which the adversary again extends by appending an edge on each side, creating an augmenting path of type 0, 1, 2, 1, 0, as shown in Fig. 5b. Since the current competitive ratio is 3/2, the algorithm applies this path. In response the adversary appends an edge at each endpoint of the type 3 edge, and at each endpoint of one of the type 1 edges, as shown in Fig. 5c. The resulting graph has a blocked augmenting path of type 0, 3, 0, and an augmenting path of type 0, 1, 0, as shown in Fig. 5c. The algorithm needs to apply the latter one as the competitive ratio is currently 5/3 > 3/2.

At this point, the adversary repeats this construction $n - 1$ times, by identifying the shaded part of Fig. 5c as the graph of Fig. 5a, and reapplying the above construction. The final graph consists of n blocked augmenting paths of type 0,3,0 and $n + 2$ edges of type 1 that belong both to the optimal and the algorithm's matchings. Hence, the competitive ratio is $(3n + 2)/(2n + 2)$, which contradicts the claimed ratio and shows a lower bound on the competitive ratio of 3/2.

Case k is even and at least 4. Fix an algorithm that claims a competitive ratio strictly smaller than $k/(k - 1)$. The adversary releases a single edge, creating an augmenting path of length 1. Whenever the algorithm applies the augmenting path,[3] the adversary extends it by appending one edge on each end, eventually creating an alternating path of type $1, 2, \ldots, k-1, k, k-1, \ldots, 2, 1$. Then the adversary appends an edge to each endpoint of the type $k - 1$ edges. The resulting graph has two augmenting paths of type $0, k - 1, 0$, see Fig. 6a. The algorithm needs to apply them as the competitive ratio is currently $(k + 2)/k$, which is strictly greater than $k/(k - 1)$ if $k \geq 4$. Each augmentation is responded, by the adversary, with an extension of the path resulting in the configuration depicted in Fig. 6b of ratio $(k + 4)/(k + 2) \geq k/(k - 1)$ where

---

[3] We can assume, without loss of generality, that this is the only viable choice for the online algorithm. This is because the only way an algorithm can transform a given matching M_1 to a matching M_2 is via a sequence which can only consist of the following: (i) augmenting paths; (ii) alternating cycles; and (iii) alternating, or even "decreasing" paths (i.e., paths such that if the algorithm applies them, then the cardinality of the matching remains the same, or decreases, respectively). This is a well-known result from matching theory. In principle an online algorithm, say A, could apply paths and cycles in cases (ii) and (iii), but such an algorithm can be converted to another algorithm A' which is at least as good as A in terms of matching size, and which maintains edges of smaller types than A.
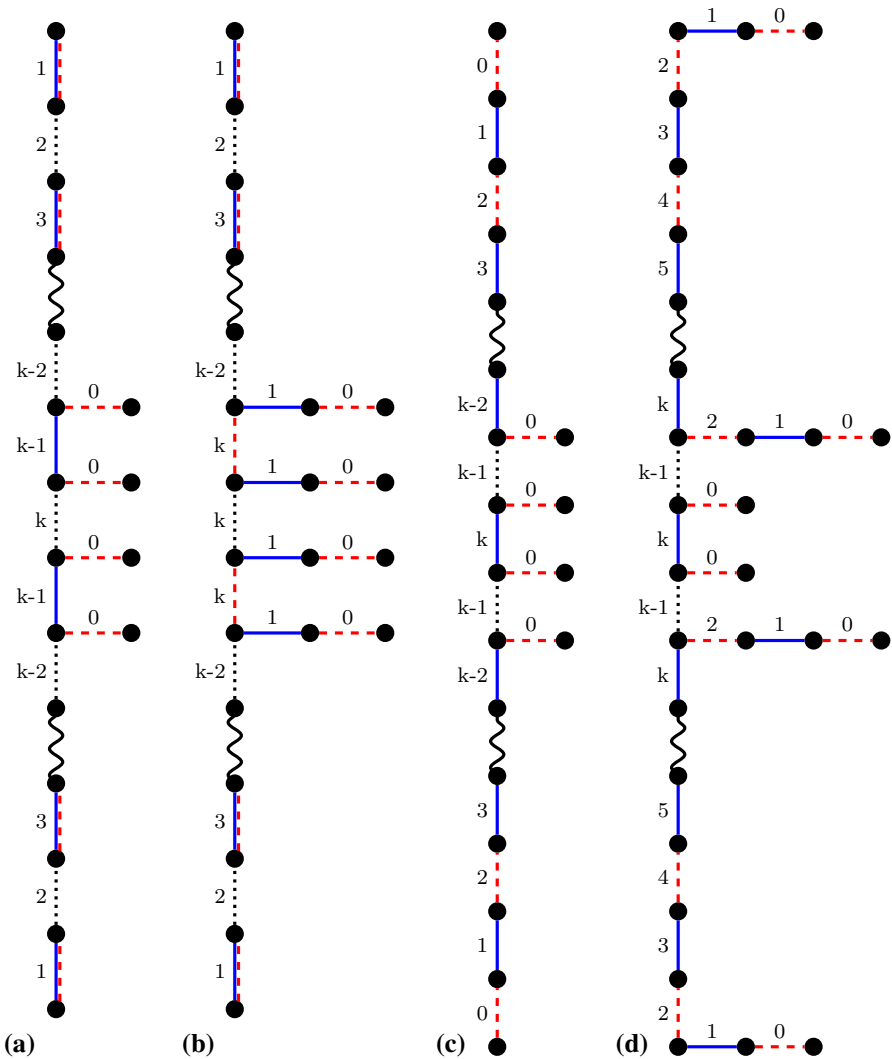
Fig. 6 Lower bound constructions for the deterministic competitive ratio. Solid/blue edges depict the algorithm's matching, dashed/red edges depict the optimal matching, dotted/black edges belong to none of the matchings and wiggled lines represent parts of the graph that are contracted for readability

all augmenting paths are blocked by type k edges. Hence the competitive ratio of the algorithm is not strictly smaller than $k/(k-1)$.

Case k is odd and at least 5. Fix an algorithm that claims a competitive ratio strictly smaller than $k/(k-1)$. The adversary proceeds as in the previous case, until the graph consists of a path of type $1, 2, \ldots, k-1, k, k-1, \ldots, 2, 1$. This time, the adversary appends one edge to each endpoint of the type $k-1$ edges, but also appends one edge at each endpoint of the path. As a result, there are two augmenting paths of type $0, 1, 2, \ldots, k-2, 0$ and a single blocked augmenting path of type $0, k, 0$, see Fig. 6c.

Table 1 Summary of lower bounds (LB) and upper bounds on the competitive ratio for the problem, for all even k with $4 \leq k \leq 22$. The lower bounds for the (limited) arrival/departure model are discussed in Sect. 3. The analysis of L-Greedy and AMP carry through to the (limited) arrival/departure model. For $k \geq 22$, the upper bound of AMP is superior to the upper bound of L-Greedy

k	LB (arr.)	LB (arr./dep.)	*L*-Greedy	AMP
4	1.333333	1.428571	1.5	2.598076
6	1.2	1.263158	1.466667	1.869186
8	1.142857	1.179487	1.428571	1.613602
10	1.111111	1.134328	1.333333	1.480583
12	1.090909	1.106796	1.318182	1.398080
14	1.076923	1.088435	1.307692	1.341500
16	1.066666	1.075377	1.300000	1.300080
18	1.058823	1.065637	1.247059	1.268330
20	1.052631	1.058104	1.242105	1.243150
22	1.047619	1.052109	1.238095	1.222640

The algorithm needs to apply an augmenting path as the competitive ratio is currently $(k + 3)/k$, which is strictly greater than $k/(k - 1)$ for $k \geq 5$. The adversary responds each augmentation of a path by appending an edge on both ends of this path. At this moment, the ratio decreased slightly, but still exceeds the claimed ratio, forcing the algorithm to continue augmenting. Eventually this leads to a configuration formed by two blocked augmenting paths of type $0, 1, 2, \ldots, k, 2, 1, 0$ and a blocked augmenting path of type $0, k, 0$, see Fig. 6d. The competitive ratio of this configuration is $(k + 7)/(k + 4)$ which exceeds $k/(k - 1)$ for $k \geq 5$.   □

2.5 Comparing the algorithms *L*-Greedy and AMP

We have analyzed two deterministic online algorithms: the algorithm AMP, which has competitive ratio $1 + O\left(\log k/k\right)$, and the algorithm L-Greedy, which has competitive ratio $1 + \Theta(1/\sqrt{k})$. Since the analysis of L-Greedy is tight, it follows that AMP is asymptotically (i.e., for large k) superior to L-Greedy. However, for small values of k, namely $k \leq 20$, we observe that L-Greedy performs better, in comparison to the performance bound we have shown for AMP. These findings are summarized in Table 1 and Fig. 2 (Sect. 1.4).

3 Online matching in the edge arrival/departure model

In this section we consider the online matching problem with k edge-recourse in the setting in which edges may arrive but also *depart* online. More precisely, a request sequence for this problem is of the form $(p_i, e_i)_{i \geq 1}$, namely the i-th request consists of an edge e_i and its *mode* $p_i \in \{\texttt{arrive}, \texttt{depart}\}$. If $p_i = \texttt{arrive}$ then the edge e_i becomes available; this corresponds to the arrival setting studied in Sect. 2.

If $p_i = \texttt{depart}$, then e_i is removed from the graph, and can be used by neither the online algorithm or the optimal offline algorithm. We emphasize that in this model, an edge of the form (u, v) may arrive and depart several times in the course of serving a request sequence, but every time it arrives it is considered as a "fresh" edge. As a consequence, upon each arrival, such an edge is assigned type 0. Moreover, a departing edge ceases to exist in any matchings.

As explained in the Introduction, we will further distinguish between two models. In the *limited departure model*, an edge cannot depart while it is being used in the matching of the online algorithm, whereas in the stronger *full departure model* any edge can depart.

It turns out that the full departure model is quite restrictive. This is because it is possible for the adversary to force an online algorithm to augment some augmenting path and then to remove one of the edges in its matching. Eventually the algorithm can end up with blocked edges (type k), without having the chance to augment its matching. This intuition is formalized in the following lemma.

Lemma 3 *The competitive ratio of online matching with k edge-recourse in the full departure model is* 2.

Proof To show an upper bound of 2, consider an algorithm which adds to its matching any edge whose endpoints are unmatched. The edge types are either 0 or 1, and therefore the algorithm is not sensitive to the given edge budget. The matching produced by the algorithm is (inclusion wise) a maximal matching, and it is well known that its size is at least 1/2 the size of the maximum matching.

To show a lower bound of 2, consider a graph consisting of vertices 1, 2, 3, 4, with the edge $(2, 3)$ arriving at the beginning. Then the edges $(1, 2)$, $(3, 4)$ arrive and depart repeatedly, alternating between two configurations. When the graph consists of the single edge $(2, 3)$, the algorithm needs to include it in its matching. When the graph consists of the path $(1, 2, 3, 4)$, the algorithm needs to apply this augmenting path if it claims to have a competitive ratio strictly lower than 2. As a result, the type of the edge $(2, 3)$ keeps increasing, and when it reaches k, the algorithm cannot augment the matching anymore. Thus, for even k, the algorithm has a matching of size 0, while the optimal matching consists of the edge $(2, 3)$. Similarly, for odd k, the algorithm has a matching of size 1, while the optimal matching has size 2. Hence, no algorithm can achieve a competitive ratio strictly lower than 2. □

Since the full departure model is very restrictive for the algorithm, as shown in Lemma 3, we will concentrate on the limited departure model, as defined in the introduction. For this model, we observe that the algorithms L-Greedy and AMP have the same performance guarantee as in the edge arrival model. This is because the analysis of L-Greedy uses weights on vertices which are not affected by edge departures, and the analysis of AMP is based on an upper bound over the number of type k edges, which still holds under edge departures. We thus focus on obtaining stronger lower bounds in this model (also included in Table 1). We begin by observing that the bound of 3/2 of the competitive ratio in the edge arrival model for $k \in \{2, 3\}$ still holds for the limited departure model, in which the adversary is stronger. Hence, the smallest interesting value for k in this model is $k = 4$, for which we provide the following specific

lower bound. The proof will provide some intuition about the adversarial argument for general k, which is shown in Theorem 5.

Theorem 4 *The competitive ratio of online matching with k edge-recourse in the* limited departure model *is at least* $10/7$ *for $k = 4$.*

Proof We will prove the theorem by applying a game between the online algorithm and the adversary. In particular, the adversary will enforce arrivals and departures of edges in such a way that, at every moment in time, the symmetric difference between the matching produced by the algorithm and the optimal matching consists only of augmenting paths. In particular, this symmetric difference will have no alternating cycles or alternating paths of even length.

Any such augmenting path can be represented as a *string* of integers in $\{0, 1, \ldots, k\}$, which is precisely the type of this path. Note that for an augmenting path, this string begins and ends with 0. We can thus think of the above-defined symmetric difference as a collection of strings, which in turn allows us to define the game between the algorithm and the adversary over this collection of strings as opposed to defining it over the actual graph.

Whenever the algorithm applies an augmenting path, this translates into the increment of all edge types of the corresponding string, for example the string 01210 becomes 12321. The adversary will respond to this augmentation by a combination of the following three possible types of *operations*.

– The adversary may *append* 0's to both ends of a string, for example $101 \rightarrow 01010$. To do so, the adversary releases two new edges (of type 0). Each edge is incident with only one endpoint of the path described by the string, and is not incident with any other edge in the current graph.
– The adversary may *split* the string into smaller strings, for example $12321 \rightarrow \{123, 1\}$ or $12321 \rightarrow \{1, 3, 1\}$. This can be done via the departure of certain edges which are not in the algorithm's matching (of even type). For instance, the operation $12321 \rightarrow \{123, 1\}$ can be done by having one edge of type 2 depart, and the operation $12321 \rightarrow \{1, 3, 1\}$ can be done by having both edges of type 2 depart.
– The adversary may *merge* certain strings, for example $\{1, 1\} \rightarrow 101$. This can be done via the arrival of a new edge (of type 0).

The game between the algorithm and the adversary   The main idea behind the adversarial construction is as follows. We suppose that the online algorithm has competitive ratio at most $(10 - \epsilon)/7$ for arbitrarily small $\epsilon > 0$. We will then show that the adversary can eventually force the algorithm to a competitive ratio at least $10/7$, thus leading to a contradiction.

The game begins with the adversary presenting the string 0, which the algorithm has to augment, resulting in a single string 010. From this point onwards, the game proceeds in three *phases*, which are depicted in Fig. 7. In each phase, a sequence of algorithm/adversary *actions* takes place. Each action is of the following form: The algorithm chooses some string s to augment, which results in a string s'. Then the adversary will perform a sequence of the above defined operations on s', which will
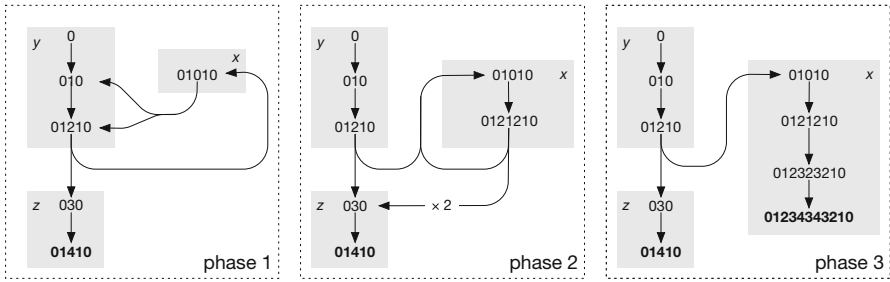
Fig. 7 An illustration of the three phases in the game between the algorithm and the adversary, for the proof of Theorem 4. Blocked strings are depicted in bold face. The arcs illustrate the actions of the adversary, after an augmentation by the algorithm. For example, if the algorithm augments the string 01010 in phase 1, then the adversary replaces the resulting string by the strings 010 and 01210, whereas in phases 2 and 3 it replaces it with the string 0121210. The numbers x, y, z count the number of strings which belong to the corresponding shown boxes

result in either a single new string (say $\bar{s}$), or to several new strings, say $\bar{s}_1, \bar{s}_2, \ldots \bar{s}_k$ (in our construction, it will be that $k \in [1, 3]$). This adversarial action is depicted by means of an arrow from s to each of the $\bar{s}_1, \bar{s}_2, \ldots \bar{s}_k$ in Fig. 7.

As an example, in phase 1, if the algorithm augments string $s = 010$ (thus obtaining string $s' = 121$), the adversary appends two zeros at both ends of s', which results in the string $\bar{s}_1 = 01210$. If the algorithm augments string $s = 01210$, thus obtaining string $s' = 12321$, then the adversary first splits s' to two strings 1, 3 and 1, which he then transforms into the strings $\bar{s}_1 = 01010$ and $\bar{s}_2 = 030$, by merging them and appending zeros. If the algorithm augments string $s = 030$, thus obtaining string $s' = 141$, then the adversary appends two zeros at both ends of s', which results in the string $\bar{s}_1 = 01410$. Last, if the algorithm augments string $s = 01010$ (thus obtaining $s' = 12121$) then the adversary first splits s' to two strings 1 and 121, then appends two zeros to the end of each string. This results in two strings $\bar{s}_1 = 010$ and $\bar{s}_2 = 01210$. The above are all possible actions that can occur in phase 1. Actions for phases 2 and 3 are similar and defined by the graphs in Fig. 7.

We also need to explain how the game *transitions* between phases; i.e., under which conditions the game moves from phase i to phase $i + 1$. To this end, we define some variables which count the numbers of some specific strings. More precisely:

– x denotes the number of strings 01010, 0121210, 012323210 or 01234343210; such strings have local ratio 3/2, 4/3, 5/4 and 6/5, respectively.
– y denotes the number of strings 0, 010 or 01210; such strings have local ratio ∞, 2 and 3/2, respectively.
– z denotes the number of strings 030 or 01410; such strings have local ratio 2 and 3/2, respectively.

In Fig. 7 we use "boxes" to illustrate this grouping of strings.

In particular, we will call strings 030 or 01410 *bad strings*. This is motivated by the observation that 01410 is *blocked* and has large local ratio equal to 3/2; note that the algorithm cannot augment such a string. Moreover, the string 030 has local ratio 2, and immediately after an augmentation it becomes 01410.

The goal of the adversary is to reach a configuration with only blocked strings 01410 and 01234343210 with a maximum proportion of bad strings 01410, since this maximizes the competitive ratio. It is relatively easy for the adversary to generate bad strings, but this comes at the expense of generating strings 01010. The adversary's goal is to minimize the proportion of these strings, and this is done through three different phases. At a high level, the objective of phase 1 is to create a large number of bad strings. This is is also the objective of phase 2, but with a more *efficient* generation of bad strings, in the sense that fewer 01010 strings are generated per bad string. The objective of phase 3 is simply to bring the game in a configuration with only blocked strings.

The game starts with the adversary entering phase 1 and generating the single string 0. In phase 1, immediately after each action of the adversary, the algorithm has competitive ratio at least $3/2$, meaning that it is forced to augment strings, since $3/2 > (10 - \epsilon)/7$, which is the competitive ratio claimed by the algorithm. This is because in this phase all strings have local ratio at least $3/2$. Throughout phase 1 we have the invariant

$$2x + y = z + 1, \tag{Inv 1}$$

which can be verified by inspecting each possible action of phase 1. For example the augmentation of 01010 decrements x and increments y by 2. Eventually the inequality $7z + 3 > 2/\epsilon$ will come to hold, simply because after at most $x + 1$ augmentations, the counter z increases strictly. At that moment, the adversary moves to phase 2.

Throughout phase 2 we have the invariants

$$7z + 3 > 2/\epsilon \tag{Inv 2.1}$$

and

$$2x + y \leq z + 1. \tag{Inv 2.2}$$

(Inv 2.1) holds because the left-hand side will not decrease throughout the phase. To show (Inv 2.2), we first observe that by invariant (Inv 1) phase 2 starts with equality, and the actions of phase 2 preserve the inequality (Inv 2.2), which can be easily shown by inspecting each possible action during phase 2.

Phase 2 continues for as long as $z < 8(x + y)$, and phase 3 begins at the point in which $z \geq 8(x + y)$. This condition will eventually be reached, because the quantity $x + y$ is invariant during phase 2, whereas any sequence of at least $x + 1$ actions increases z by at least 1.

We will now argue that in phase 2, right after each action of the adversary, the competitive ratio of the algorithm is strictly greater than $(10 - \epsilon)/7$, which implies that the algorithm must, in turn, respond with an augmentation to every action of the adversary in phase 2, since we assumed that the algorithm is $(10 - \epsilon)/7$-competitive. To this end, we observe that at each point in phase 2, the algorithm maintains certain types of strings whose local ratio we lower bounded above. In particular, there are $y + z$

strings of local ratio at least $3/2$ and x strings of local ratio at least $4/3$. Therefore, a lower bound to the competitive ratio during phase 1, can be stated as follows, where we will make use of the property

$$f(x, y) := \frac{ax + cy}{bx + dy} \quad \text{is decreasing on } x \text{ and increasing on } y \text{ if} \qquad \text{(P1)}$$

$$a, b, c, d > 0 \text{ and } \frac{a}{b} \leq \frac{c}{d}.$$

We have

$$
\begin{aligned}
\frac{4x + 3y + 3z}{3x + 2y + 2z} &= \frac{8x + 6y + 6z}{6x + 4y + 4z} \qquad \text{(from P1 and Inv 2.2)} \\
&= \frac{4(2x + y) + 2y + 6z}{3(2x + y) + y + 4z} \\
&\geq \frac{4(z + 1) + 2y + 6z}{3(z + 1) + y + 4z} \\
&= \frac{2y + 10z + 4}{y + 7z + 3}
\end{aligned}
$$

$$
\begin{aligned}
&\geq \frac{10z + 4}{7z + 3} \qquad \text{(from P1)} \\
&= \frac{10}{7} - \frac{\frac{2}{7}}{7z + 3}
\end{aligned}
$$

$$
\begin{aligned}
&> \frac{10}{7} - \frac{\frac{2}{7}}{\frac{2}{\epsilon}} \qquad \text{(from Inv 1)} \\
&= \frac{10 - \epsilon}{7}.
\end{aligned}
$$

Recall that when the condition $z \geq 8(x + y)$ becomes satisfied, the game moves to the final phase, namely phase 3. Moreover, the condition $z \geq 8(x+y)$ holds throughout phase 3, since $x + y$ is invariant and z can only increase in this phase. We also obtain that

$$y + z \geq y + 8(x + y) \geq 8x,$$

which will be useful. Similar to the previous argument, we can lower bound the competitive ratio after each action of the adversary by

$$\frac{3(y+z)+6x}{2(y+z)+5x} \geq \frac{3\cdot(8x)+6x}{2\cdot(8x)+5x} \qquad \text{(From } y+z \geq 8x \text{ and P1)}$$
$$= \frac{10}{7}.$$

Therefore, the algorithm must augment after each action of the adversary, and eventually must find itself in a configuration which consists only of blocked strings (either 01410 or 01234343210). At this configuration, the algorithm cannot do any further augmentations, hence its competitive ratio is at least 10/7, a contradiction. □

We can generalize the ideas in the proof of Theorem 4 so as to obtain a non-trivial lower bound for general even $k \geq 4$ in Theorem 5. Note that since $\frac{k^2-3k+6}{k^2-4k+7} > 1 + \frac{1}{k-1}$ for all $k \geq 4$, Theorem 5 shows a stronger lower bound for even k than Theorem 3 under the limited departure model.

Theorem 5 *The competitive ratio online matching with k edge-recourse in the limited departure model is at least $\frac{k^2-3k+6}{k^2-4k+7}$, for all even $k \geq 4$.*

Proof First, we observe that for $k = 4$, the expression $\frac{k^2-3k+6}{k^2-4k+7}$ is equal to 10/7, which is precisely the value obtained in Theorem 4. Therefore, it suffices to prove the result for even $k \geq 6$.

The proof generalizes the ideas behind the proof of Theorem 4, and in particular the concept of a game between the online algorithm and the adversary. Again, we suppose that the online algorithm has competitive ratio at most $\frac{k^2-3k+6-\epsilon}{k^2-4k+7}$ for arbitrarily small $\epsilon > 0$. We will then show that the adversary can eventually force the algorithm to a competitive ratio at least $\frac{k^2-3k+6}{k^2-4k+7}$, thus leading to a contradiction.

The game begins with the adversary presenting the string 0, which the algorithm has to augment, resulting in a single string 010. From this point onwards the game proceeds in two phases, which are depicted in Fig. 8. In each phase, a sequence of algorithm/adversary actions takes place. Actions are defined to be consistent with Fig. 8 for the two phases of the game. It is worth pointing out that the game for $k \geq 6$ consists of two phases, while in contrast, the game for $k = 4$ (as described in the proof of Theorem 4) consists of three phases. The second phase for $k = 4$ is necessary for the adversary to force the algorithm to keep augmenting strings.

We also need to explain under which conditions the game transitions from phase 1 to phase 2. To this end, we define again some variables which count the numbers of certain specific strings. More precisely:

– x denotes the number of strings $0123\ldots i(i-1)i\ldots3210$ for all $1 \leq i \leq k-2$; such strings have local ratio $\frac{i+2}{i+1}$.
– a_i, where i is even in $[2, k]$, denotes the number of strings $0(i-1)0$ or $01i10$; such strings have local ratio 2 and 3/2, respectively.
– v denotes the number of strings $0123\ldots(k-3)0, 0123\ldots(k-2)10, 0123\ldots(k-1)210$ or $0123\ldots k3210$; such strings have local ratio $\frac{k/2}{k/2-1}$, $\frac{k/2+1}{k/2}$, $\frac{k/2+2}{k/2+1}$ and $\frac{k/2+3}{k/2+2}$, respectively.
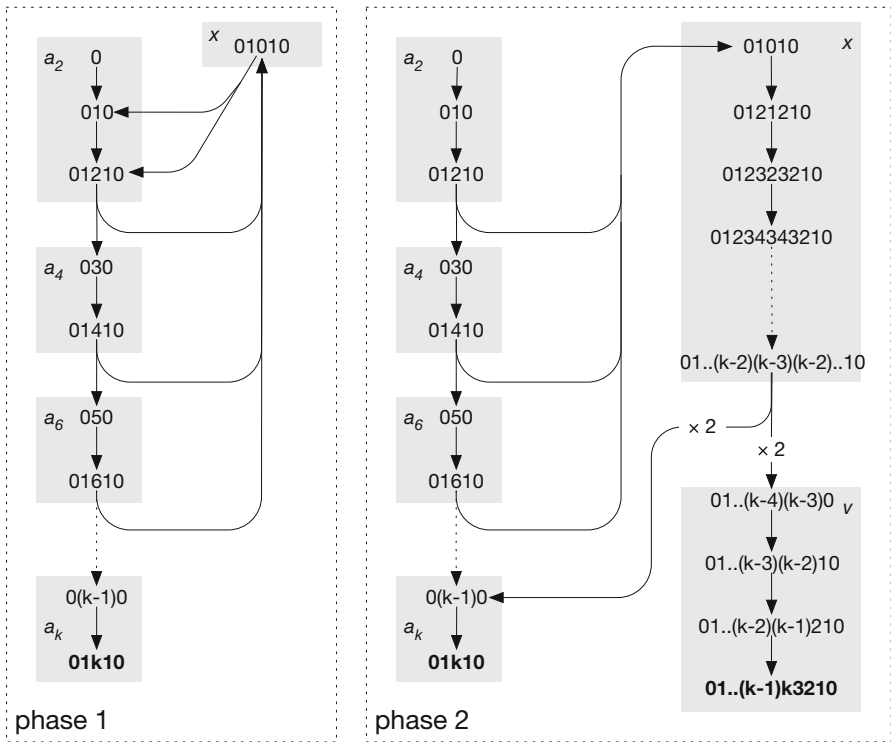
Fig. 8 The lower bound construction in the arrival/departure model. Blocked strings are depicted in bold face. The arcs illustrate the adversarial strategy. For example if the algorithm augments the string $01 \ldots (k-2)(k-3)(k-2) \ldots 10$ the adversary replaces the resulting string by two strings $0(k-1)0$ and two strings $01 \ldots (k-4)(k-3)0$. The numbers $a_2, a_4, a_6, \ldots, a_k, x, v$ count the number of strings which belong to the corresponding shown boxes

In phase 1, immediately after each action of the adversary, the algorithm has competitive ratio at least $3/2$. This is because in this phase all strings have local ratio at least $3/2$. However, after each augmentation by the algorithm the competitive ratio can be much better, and possibly smaller than $\frac{k^2-3k+6}{k^2-4k+7}$. For this reason, the adversary will move eventually the game to phase 2. In particular, phase 2 begins once the following condition is satisfied.

$$\sum_{j=1}^{k/2} (2j(k-1) - 3k + 7) \cdot a_{2j} > \frac{2(k-3)}{\epsilon} - (k-1).$$

Note that this condition will be satisfied because the coefficient $(2j(k-1) - 3k + 7)$ is non-negative for $j \geq 2$ and increasing in j for $j \geq 1$, and whenever a_{2j} decreases by one, $a_{2(j+1)}$ increases by one during the execution of phase 1. Intuitively, the objective in phase 1 is to create a large number of strings counted by $a_4, a_6, \ldots, a_k$, whereas in phase 2 the objective is to force the algorithm in a configuration with only blocked strings.

Throughout phase 2, we have the invariants

$$\sum_{j=1}^{k/2} (2j(k-1) - 3k + 7) \cdot a_{2j} > \frac{2(k-3)}{\epsilon} - (k-1), \qquad \text{(Inv 3.1)}$$

$$2x \le 1 + \sum_{j=1}^{k/2} (2j-3)a_{2j} - (k-2)v. \qquad \text{(Inv 3.2)}$$

Invariant (Inv 3.1) holds because the left-hand side will not decrease throughout the phase. To show Invariant (Inv 3.2), we first observe that at the end of phase 1, it holds that

$$2x \le 1 + \sum_{j=1}^{k/2} (2j-3)a_{2j};$$

this can be easily shown by induction on the number of actions during phase 1. Thus, at the beginning of phase 2, Invariant (Inv 3.2) holds, since $v = 0$ at that point. Again, a simple inductive argument on the number of actions throughout phase 2 can show that the invariant is maintained.

We will argue that Invariants (Inv 3.1) and (Inv 3.2) imply that throughout phase 2, the competitive ratio is strictly larger than $\frac{k^2-3k+6-\epsilon}{k^2-4k+7}$, and hence throughout phase 2 the algorithm is forced to augment strings, until all strings are blocked. This is because the competitive ratio in phase 2 can be lower bounded by

$$\frac{(\frac{k}{2}+3)v + 3\sum_{j=1}^{k/2} a_{2j} + kx}{(\frac{k}{2}+2)v + 2\sum_{j=1}^{k/2} a_{2j} + (k-1)x} \qquad \text{(From P1 and Inv 3.2)}$$

$$= \frac{(k+6)v + 6\sum_{j=1}^{k/2} a_{2j} + 2kx}{(k+4)v + 4\sum_{j=1}^{k/2} a_{2j} + 2(k-1)x}$$

$$\ge \frac{(k+6)v + 6\sum_{j=1}^{k/2} a_{2j} + k\left(1 + \sum_{j=1}^{k/2}(2j-3)a_{2j} - (k-2)v\right)}{(k+4)v + 4\sum_{j=1}^{k/2} a_{2j} + (k-1)\left(1 + \sum_{j=1}^{k/2}(2j-3)a_{2j} - (k-2)v\right)}$$

$$= \frac{(-k^2+3k+6)v + \sum_{j=1}^{k/2}(k(2j-3)+6)\,a_{2j} + k}{(-k^2+4k+2)v + \sum_{j=1}^{k/2}((k-1)(2j-3)+4)\,a_{2j} + k - 1}.$$

To complete the proof, it remains to show that

$$\frac{(-k^2 + 3k + 6)v + \sum_{j=1}^{k/2} (k(2j - 3) + 6)\, a_{2j} + k}{(-k^2 + 4k + 2)v + \sum_{j=1}^{k/2} ((k - 1)(2j - 3) + 4)\, a_{2j} + k - 1} > \frac{k^2 - 3k + 6 - \epsilon}{k^2 - 4k + 7}. \tag{13}$$

By a simple mathematical manipulation, for (13) to hold, it suffices that

$$v \cdot C_v + \sum_{j=1}^{k/2} a_{2j} \cdot C_{2j} > 2(k - 3) - (k - 1)\epsilon, \tag{14}$$

where C_v and C_{2j} are defined as

$$C_v = 3(k^2 - 7k + 10) - (k^2 - 4k - 2)\epsilon$$
$$C_{2j} = 2(k - 3)(k - 2j) + (2j(k - 1) - 3k + 7)\epsilon.$$

We observe that C_v and the first additive term in the expression of C_{2j} (i.e. $2(k-3)(k-2j)$) are non-negative for sufficiently small ϵ, $k \geq 6$ and $1 \leq j \leq k/2$. Removing these terms from the left hand side of (14), it suffices to show that

$$\sum_{j=1}^{k/2} (2j(k - 1) - 3k + 7)\epsilon \cdot a_{2j} > 2(k - 3) - (k - 1)\epsilon.$$

This inequality is precisely Invariant (Inv 3.1), which concludes the proof. $\qquad\square$

4 Conclusion

In this paper we provided improved upper and lower bounds for online maximum matching with k edge-recourse. More specifically, we analyzed two online algorithms for the edge arrival model, namely AMP and L-GREEDY which seem to be incomparable: the former is asymptotically superior, in terms of k, but the latter has a better performance analysis for small k. It would be interesting to analyze an algorithm that combines the ingredients of these two algorithms, namely, an algorithm that combines the doubling techniques with augmenting only along short paths. The difficulty in the analysis of such an algorithm lies in that reasonable charging schemes tend to have "local" properties, wheres the doubling algorithm applies a "global" criterion which does not easily translate into some structural property that can be useful in analysis.

The problems we consider remain challenging even for k as small as 4, and some gap between the upper and lower bounds remains. Bringing this gap will probably require new ideas and techniques. To this end, it is worth pointing out that the amortization arguments we used in our analysis may have connections to LP-based algorithms (since the dual of the maximum matching problem is a weighted vertex minimization problem). Thus, it would be very interesting to use a duality-based approach, such as dual fitting, towards the design and analysis of improved algorithms.

References

Avitabile T, Mathieu C, Parkinson LH (2013) Online constrained optimization with recourse. Inf Process Lett 113(3):81–86

Bernstein A, Holm J, Rotenberg E (2018) Online bipartite matching with amortized replacements. In: Proceedings of the 29th annual ACM-SIAM symposium on discrete algorithms, (SODA), pp 947–959

Borodin A, El-Yaniv R (1998) Online computation and competitive analysis. Cambridge University Press, Cambridge

Boyar J, Favrholdt LM, Kotrbčík M, Larsen KS (2017) Relaxing the irrevocability requirement for online graph algorithms. In: Proceedings of the 15th workshop on algorithms and data structures, (WADS), pp 217–228

Buchbinder N, Segev D, Tkach Y (2017) Online algorithms for maximum cardinality matching with edge arrivals. In: Proceedings of the 25th annual European symposium on algorithms (ESA), pp 22:1–22:14

Chiplunkar A, Tirodkar S, Vishwanathan S (2015) On randomized algorithms for matching in the online preemptive model. In: Proceedings of the 23rd annual european symposium on algorithms (ESA), pp 325–336

Epstein L, Levin A, Segev D, Weimann O (2013) Improved bounds for online preemptive matching. In: Proceedings of the 30th international symposium on theoretical aspects of computer science (STACS), pp 389–399

Gu A, Gupta A, Kumar A (2016) The power of deferral: maintaining a constant-competitive steiner tree online. SIAM J Comput 45(1):1–28

Gupta A, Kumar A (2014) Online Steiner tree with deletions. In: Proceedings of the 25th annual ACM-SIAM symposium on discrete algorithms, (SODA), Portland, Oregon, USA, pp 455–467

Gupta A, Kumar A, Stein C (2014) Maintaining assignments online: matching, scheduling, and flows. In: Proceedings of the 25th snnual ACM-SIAM symposium on discrete algorithms (SODA), pp 468–479

Han X, Makino K (2009) Online minimization knapsack problem. In: Proceedings of the 7th international workshop on approximation and online algorithms (WAOA), pp 182–193

Iwama K, Taketomi S (2002) Removable online knapsack problems. In: Proceedings of the 29th international colloquium on automata, languages and programming (ICALP), pp 293–305

Karp RM, Vazirani UV, Vazirani VV (1990) An optimal algorithm for on-line bipartite matching. In: Proceedings of the 22nd annual ACM symposium on theory of computing (STOC), pp 352–358. ACM

Love E (1980) Some logarithmic inequalities. Math Gaz 64(427):55–57

McGregor A (2005) Finding graph matchings in data streams. In: Approximation, randomization and combinatorial optimization, algorithms and techniques (APPROX-RANDOM), pp 170–181

Megow N, Skutella M, Verschae J, Wiese A (2016) The power of recourse for online MST and TSP. SIAM J Comput 45(3):859–880

Mehta A (2013) Online matching and ad allocation. Found Trends Theor Comput Sci 8(4):265–368

Siegel C (1988) Topics in complex function theory: elliptic functions and uniformization theory, vol 1. Wiley, New York

Varadaraja AB (2011) Buyback problem-approximate matroid intersection with cancellation costs. In: Proceedings of the 38th international colloquium on automata, languages, and programming (ICALP), pp 379–390

Affiliations

Spyros Angelopoulos[1] · **Christoph Dürr[1]** · **Shendan Jin[1]**

Christoph Dürr
christoph.durr@lip6.fr

Shendan Jin
shendan.jin@lip6.fr

[1]   Laboratoire d'informatique de Paris 6, LIP6, CNRS, Sorbonne Université, 75252 Paris, France