



Approximating Multistage Matching Problems

Markus Chimani¹ · Niklas Troost¹ · Tilo Wiedera¹

Received: 31 August 2021 / Accepted: 7 February 2022 / Published online: 23 March 2022
© The Author(s) 2022

Abstract

In *multistage perfect matching* problems, we are given a sequence of graphs on the same vertex set and are asked to find a sequence of perfect matchings, corresponding to the sequence of graphs, such that consecutive matchings are as similar as possible. More precisely, we aim to maximize the intersections, or minimize the unions between consecutive matchings. We show that these problems are NP-hard even in very restricted scenarios. As our main contribution, we present the first non-trivial approximation algorithms for these problems: On the one hand, we devise a tight approximation on graph sequences of length two (2-stage graphs). On the other hand, we propose several general methods to deduce multistage approximations from black-box approximations on 2-stage graphs.

Keywords Temporal graphs · Approximation algorithms · Perfect matchings

1 Introduction

The study of graphs that evolve over time emerges naturally in several applications. As such, it is a well-known subject in graph theory [1–11, 13, 14, 16–18, 21, 24]. While there are many possible approaches to model these problems (see the discussion of related work), the paradigm of *multistage graphs* has attracted quite some attention in recent years [2–5, 10, 14, 16, 17]. In this setting, we are given a sequence of separate, but related graphs (*stages*). A typical goal is to find a sequence of solutions for each individual graph such that the change in the solutions between consecutive graphs is minimized. Since multistage graph problems usually turn out to be NP-hard, one

✉ Niklas Troost
niklas.troost@uos.de

Markus Chimani
markus.chimani@uos.de

Tilo Wiedera
tilo.wiedera@uos.de

¹ Theoretical Computer Science, Osnabrück University, Osnabrück, Germany

often resorts to FPT- or approximation algorithms. To the best of our knowledge, all approximation results in this setting discuss combined objective functions that reflect a trade-off between the quality of each individual solution and the cost of the change over time (see, e.g., [2, 16]). However, this is a drawback if one requires each stage's solution to attain a certain quality guarantee (e.g., optimality). Trying to ensure this by adjusting the trade-off weights in the above approximation algorithms leads to approximation ratios that no longer effectively bound the cost of change. Here, we discuss a multistage graph problem where each individual solution is necessarily optimal, but we can still obtain an approximation ratio on the cost of the change over time.

A classical example are *multistage matching* problems, i.e., natural multistage generalizations of traditional matching problems (e.g., perfect matching, maximum weight matching, etc.). This is particularly interesting as optimality for a single stage would be obtainable in polynomial time, but all known multistage variants are NP-hard already for two stages. There are several known approximation algorithms for multistage matching problems [2]; however, they all follow the trade-off paradigm.

In this paper, we are concerned with maintaining a perfect matching on a multistage graph, such that the changes between consecutive matchings are minimized. After showcasing the complexities of our problems (Sect. 2), we will devise efficient approximation algorithms (Sect. 3).

1.1 Definitions and Preliminaries

Let $G = (V, E)$ be an undirected graph. For a set $W \subseteq V$ of vertices, let $\delta(W) := \{uv \in E \mid u \in W, v \in V \setminus W\}$ denote the set of its cut edges. For a singleton $\{v\}$, we may write $\delta(v)$ instead of $\delta(\{v\})$. A set $M \subseteq E$ of edges is a *matching* if every vertex is incident to at most one edge of M ; it is a *perfect matching* if $|\delta(v) \cap M| = 1$ for every $v \in V$. A k -*cycle* (k -*path*) is a cycle (path, respectively) consisting of exactly k edges. The parity of a k -cycle is the parity of k . For a set $F \subseteq E$ of edges, let $V(F) := \{v \in V \mid \delta(v) \cap F \neq \emptyset\}$ denote its incident vertices.

For $x \in \mathbb{N}$, we define $[x] := \{1, \dots, x\}$ and $\llbracket x \rrbracket := \{0\} \cup [x]$. A *temporal graph* (or τ -*stage graph*) is a tuple $\mathcal{G} = (V, E_1, \dots, E_\tau)$ consisting of a vertex set V and multiple edge sets E_i , one for each $i \in [\tau]$. The graph $G_i := (V(E_i), E_i)$ is the i th stage of \mathcal{G} . We define $n_i := |V(E_i)|$, and $n := |V|$. A temporal graph is *spanning* if $V(E_i) = V$ for each $i \in [\tau]$.

Let $\mu := \max_{i \in [\tau-1]} |E_i \cap E_{i+1}|$ denote the maximum number of edges that are common between two adjacent stages. Let $E_\cap := \bigcap_{i \in [\tau]} E_i$ and $E_\cup := \bigcup_{i \in [\tau]} E_i$. The graph $G_\cup := (V(E_\cup), E_\cup)$ is the *union graph* of \mathcal{G} . A *multistage perfect matching* in \mathcal{G} is a sequence $\mathcal{M} := (M_i)_{i \in [\tau]}$ such that for each $i \in [\tau]$, M_i is a perfect matching in G_i .

All problems considered in this paper (MIM, MUM, Min-MPM, Max-MPM; see below) are of the following form: Given a temporal graph \mathcal{G} , we ask for a multistage perfect matching \mathcal{M} optimizing some objective function. In their respective decision variants, the input furthermore consists of some value κ and we ask whether there is

an \mathcal{M} with objective value at most (minimization problems) or at least (maximization problems) κ .

Definition 1 (MIM and τ -IM) Given a temporal graph \mathcal{G} , the *multistage intersection matching problem* (MIM) asks for a multistage perfect matching \mathcal{M} of \mathcal{G} with maximum profit $p(\mathcal{M}) := \sum_{i \in [\tau-1]} |M_i \cap M_{i+1}|$. For fixed τ , we denote the problem by τ -IM.

We also consider the natural inverse objective, i.e., minimizing the unions. While the problems differ in the precise objective function, an optimal solution of MIM is optimal for MUM as well, and vice versa.

Definition 2 (MUM and τ -UM) Given a temporal graph \mathcal{G} , the *multistage union matching problem* (MUM) asks for a multistage perfect matching \mathcal{M} of \mathcal{G} with minimum cost $c(\mathcal{M}) := \sum_{i \in [\tau-1]} |M_i \cup M_{i+1}|$. For fixed τ , we denote the problem by τ -UM.

Consider either MIM or MUM. Given a temporal graph \mathcal{G} , we denote with opt the optimal solution value and with apx the objective value achieved by a given algorithm with input \mathcal{G} . The *approximation ratio* of an approximation algorithm for MIM (MUM) is the infimum (supremum, respectively) of apx/opt over all instances.

1.2 Related Work

The classical *dynamic graph* setting often considers small modifications, e.g., single edge insertions/deletions [13, 24]. There, one is given a graph with a sequence of modifications and asked for a feasible solution after each modification. A natural approach to tackle matchings in such graphs is to make local changes to the previous solutions [7–9, 23].

A more general way of modeling changes is that of *temporal graphs*, introduced by Kempe et al. [18] and used herein. Typically, each vertex and edge is assigned a set of time intervals that specify when it is present. This allows an arbitrary number of changes to occur at the same time. Algorithms for this setting usually require a more global perspective and many approaches do not rely solely on local changes. In fact, many temporal (matching) problems turn out to be hard, even w.r.t. approximation and fixed-parameter-tractability [1, 6, 11, 20, 21].

One particular flavor of temporal graph problems is concerned with obtaining a sequence of solutions—one for each stage—while optimizing a global quantity. These problems are often referred to as *multistage problems* and gained much attention in recent years [2–5, 14, 16, 17], including in the realm of matchings: e.g., the authors of [17] show W[1]-hardness for finding the largest edge set that induces a matching in each stage.

In the literature we find the problem Max-MPM, where the graph is augmented with time-dependent edge weights, and we want to maximize the value of each individual perfect matching (subject to the given edge costs) *plus* the total profit [2]. MIM is the special case where all edge costs are zero, i.e., we only care about the multistage properties of the solution, as long as each stage is perfectly matched. There is also the

inverse optimization problem Min-MPM, where we minimize the value of each perfect matching *plus* the number of matching edges newly introduced in each stage. We have APX-hardness for Max-MPM and Min-MPM [2, 16] (for Min-MPM one may assume a complete graph at each stage, possibly including edges of infinite weight). The latter remains APX-hard even for spanning 2-stage graphs with bipartite union graph and no edge weights (i.e., we only minimize the number of edge swaps) [2]. For uniform edge weights 0, the objective of Min-MPM is to minimize $\sum_{i \in [\tau-1]} |M_{i+1} \setminus M_i|$; similar but slightly different to MUM (equal up to additive $\sum_{i \in [\tau-1]} n_i/2$). For Min-MPM on metric spanning 2- or 3-stage graphs, the authors of [2] show 3-approximations. They also propose a $(1/2)$ -approximation for Max-MPM on spanning temporal graphs with any number of stages, which is unfortunately wrong (see “Appendix A” for a detailed discussion).

When restricting Max-MPM and Min-MPM to uniform edge weights 0, optimal solutions for MIM, MUM, Max-MPM, and Min-MPM are identical; thus MIM and MUM are NP-hard. However, the APX-hardness of Min-MPM does not imply APX-hardness of MUM as the objective functions slightly differ. Furthermore, the APX-hardness reduction to Max-MPM inherently requires non-uniform edge weights and does not translate to MIM. To the best of our knowledge, there are no non-trivial approximation algorithms for any of these problems on more than three stages.

1.3 Our Contribution

We start with showing in Sect. 2 that the problems are NP-hard even in much more restricted scenarios than previously known, and that (a lower bound for) the integrality gap of the natural linear program for 2-IM is close to the approximation ratio we will subsequently devise. This hints that stronger approximation ratios may be hard to obtain, at least using LP techniques.

As our main contribution, we propose several approximation algorithms for the multistage problems MIM and MUM, as well as for their stage-restricted variants, see Fig. 1. In particular, in Sect. 3.1, we show a $(1/\sqrt{2\mu})$ -approximation for 2-IM and that this analysis is tight. Then, in Sect. 3.2, we show that any approximation of 2-IM can be used to derive two different approximation algorithms for MIM, whose approximation ratios are a priori incomparable. In Sect. 3.3, we further show how to use all these algorithms to approximate MUM (and 2-UM). We also observe that it is infeasible to use an arbitrary MUM algorithm to approximate MIM. In particular, we propose the seemingly first approximation algorithms for MIM and MUM on arbitrarily many stages. We stress that our goal is to always guarantee a perfect matching in each stage; the approximation ratio deals purely with optimizing the transition costs. Recall that approximation algorithms optimizing a weighted sum between intra- and interstage costs cannot guarantee such solutions in general.

1.4 Preprocessing and Observations

Given a graph $G = (V, E)$, a single edge e is *allowed* if there exists a perfect matching M in G with $e \in M$ and *forbidden* otherwise. A graph is *matching-covered* if




Fig. 1 Relations of our approximation results. An arc from problem A to B labeled $f(\alpha)$ denotes the existence of an $f(\alpha)$ -approximation for B , given an α -approximation for A . In Corollary 16, α has to be constant. In Corollary 15, $\alpha(\cdot)$ is a function of μ . The ratio of 2-IM is by Theorem 8; combining this with Theorem 17 yields the ratio for 2-UM

all its edges are allowed (see [19] for a concise characterization of matching-covered graphs). Forbidden edges can easily be found in polynomial time; see e.g. [22] for an efficient algorithm. A simple preprocessing for MIM and MUM is to remove the forbidden edges in each stage, as they will never be part of a multistage matching. Thereby, we obtain an equivalent *reduced* temporal graph, i.e., a temporal graph whose stages are matching-covered. If any stage in the reduced temporal graph contains no edges (but vertices), the instance is infeasible. In the following, we thus assume w.l.o.g. that the given temporal graph is reduced and *feasible*, i.e., in each stage there exists some perfect matching.

Observation 3 *Let \mathcal{G} be a reduced 2-stage graph. For any $e \in E_\cap$, there is a perfect matching in each stage that includes e . Thus, there is a multistage perfect matching with profit at least 1 if $E_\cap \neq \emptyset$.*

Observation 4 *For any multistage perfect matching $(M_i)_{i \in [\tau]}$, it holds for each $i \in [\tau - 1]$ that $\max(n_i/2, n_{i+1}/2) \leq |M_i \cup M_{i+1}| = c(M_i, M_{i+1}) \leq 2 \max(n_i/2, n_{i+1}/2)$. Thus, computing any multistage perfect matching is an immediate 2-approximation for MUM.*

Observation 5 *Consider the following algorithm: Enumerate every possible sequence $(F_i)_{i \in [\tau-1]}$ such that $F_i \subseteq E_i \cap E_{i+1}$ for each $i \in [\tau - 1]$; then check for each $i \in [\tau]$ whether there is a perfect matching M_i in G_i such that $F_{i-1} \cup F_i \subseteq M_i$, where $F_0 = F_\tau = \emptyset$. Thus, MIM and MUM are in FPT w.r.t. parameter $\sum_{i \in [\tau-1]} |E_i \cap E_{i+1}|$ (or similarly $\tau \cdot \mu$).*

2 Setting the Ground

Before we present our main contribution, the approximation algorithms, we motivate the intrinsic complexities of the considered problems. On the one hand, we show that the problem is already hard in very restricted cases. On the other hand, we show that natural linear programming methods cannot yield a constant-factor approximation for 2-IM.




Fig. 2 Theorem 6: E_1 is curvy blue, E_2 is straight red. Marked vertices are encircled (Color figure online)

2.1 NP-Hardness

While it is known that 2-IM is NP-hard in general, we show that 2-IM is already NP-hard in the seemingly simple case where each vertex has only degree 2 in both stages. It immediately follows that the decision variants of MIM, 2-UM, MUM, Min-MPM, and Max-MPM remain NP-hard as well, even if restricted to this set of temporal graphs.

Theorem 6 *Deciding 2-IM is NP-hard on spanning temporal graphs whose union graph is bipartite, even if both stages consist only of disjoint even cycles and E_{\cap} is a collection of disjoint 2-paths.*

Proof We will perform a reduction from MaxCut [15] to 2-IM. In MaxCut, one is given an undirected graph $G = (V, E)$, a natural number k and the question is to decide whether there is an $S \subseteq V$ such that $|\delta(S)| \geq k$. In the first stage, we will construct an even cycle for each vertex and each edge of the original graph and in the second stage we will create an even cycle for each incidence between an edge and a vertex (see Fig. 2). A perfect matching in the first stage will correspond to a vertex selection and a perfect matching in the second stage will allow us to count the edges that are incident to exactly one selected vertex.

Given an instance $\mathcal{I} := (G = (V, E), k)$ of MaxCut, we construct an instance $\mathcal{J} := (\mathcal{G}, \kappa)$ of 2-IM. Set $\kappa := 3|E| + k$. We start with an empty 2-stage graph $\mathcal{G} := (V', E_1, E_2)$.

Let $I := \{(v, e) \mid v \in V, e \in \delta(v)\}$ be the set of incidences. For each $(v, e) \in I$, we add two new disjoint 2-paths to $E_1 \cap E_2$ and call them X_v^e and Y_v^e . Mark one edge of each X_v^e as x_v^e and one edge of each Y_v^e as y_v^e . We will refer to the endpoint of X_v^e (Y_v^e) incident to x_v^e (y_v^e) as the *marked endpoint* of X_v^e (respectively Y_v^e).

In G_1 , for each $e = vw \in E$, we create a 6-cycle through Y_v^e and Y_w^e by adding an edge between the marked endpoint of one path and the unmarked endpoint of the other, and vice versa. Furthermore, for each $v \in V$, we create a cycle of length $4|\delta(v)|$ through the paths X_v^e for $e \in \delta(v)$ as follows: For each edge $e \in \delta(v)$ and its successor edge f , according to some arbitrary cyclic order of $\delta(v)$, connect the marked endpoint of X_v^e to the unmarked endpoint of X_v^f by a 2-path with a new inner vertex.

In G_2 , for each $(v, e) \in I$, we generate a 6-cycle through X_v^e and Y_v^e by adding an edge between the marked and an edge between the unmarked endpoints of the 2-paths,

respectively. G_1 consists of $|V| + |E|$ and G_2 of $2|E|$ disjoint even cycles, thus \mathcal{G} is reduced.

Claim \mathcal{J} is a yes-instance if and only if \mathcal{I} is a yes-instance.

Proof of Claim Since both stages of \mathcal{G} consist only of pairwise disjoint even cycles and there are only two perfect matchings in an even cycle, a perfect matching in a stage is determined by choosing one edge in each cycle. For $e = vw \in E$, let $X^e := E(X_v^e) \cup E(X_w^e)$ and $Y^e := E(Y_v^e) \cup E(Y_w^e)$. Observe that for any multistage perfect matching (M_1, M_2) , $M_1 \cap M_2 \subseteq E_\cap = \bigoplus_{e \in E} X^e \cup Y^e$.

“ \Leftarrow ” Suppose there is an $S \subseteq V$, such that $|\delta(S)| \geq k$. For each $(v, e) \in I$, add x_v^e to both M_1 and M_2 if $v \in S$. Otherwise, add the unmarked edge of X_v^e to M_1 and M_2 . This uniquely determines a perfect matching M_2 in G_2 , where $y_v^e \in M_2 \iff x_v^e \in M_2 \iff v \in S$. For each $e = vw \in \delta(S)$ with $v \in S$ and $w \notin S$, add y_v^e to M_1 ; thus $y_w^e \notin M_1$ and $y_v^e \in M_1 \cap M_2$. For $e = vw \notin \delta(S)$, add either y_v^e or y_w^e to M_1 (chosen arbitrarily). This determines a 2-stage perfect matching (M_1, M_2) .

Consider some edge $e = vw \in E$. The intersection $M_1 \cap M_2$ contains two edges of X^e . If $e \in \delta(S)$, it also contains two edges of Y^e , one marked and one unmarked. If $e \notin \delta(S)$, $M_1 \cap M_2$ contains exactly one edge of Y^e . Thus, $|M_1 \cap M_2| = |\bigoplus_{e \in E} M_1 \cap M_2 \cap X^e| + |\bigoplus_{e \in E} M_1 \cap M_2 \cap Y^e| = 3|E| + |\delta(S)|$.

“ \Rightarrow ” Let (M_1, M_2) be a multistage perfect matching in \mathcal{G} with $|M_1 \cap M_2| \geq 3|E| + k$. Observe that $|M_1 \cap M_2| = \sum_{e \in E} m_e$ with $m_e := |M_1 \cap M_2 \cap (X^e \cup Y^e)| \leq 4$ for each $e \in E$. Thus, by pigeonhole principle, there are at least k edges with $m_e = 4$.

M_1 yields a selection $S \subseteq V$: Select $v \in V$ if and only if the set $X(v) := \{x_v^e \mid e \in \delta(v)\}$ is contained in M_1 . In a perfect matching in G_1 , either all or none of the edges in $X(v)$ are matched simultaneously. It can be seen that $m_e = 4$ if and only if $e \in \delta(S)$, thus $|\delta(S)| \geq k$. □

The cycles of length $4|\delta(v)|$ may have introduced an even number of vertices $W \subseteq V$ that are isolated in G_2 . To make \mathcal{G} spanning, we add to E_2 an even cycle on W . This neither interferes with E_\cap nor the profit p , since W is an independent set in the first stage G_1 . □

2.2 Linear Programming Gap

Linear Programs (LPs)—as relaxations of integer linear programs (ILPs)—are often used to provide dual bounds in the approximation context. Here, we consider the natural LP-formulation of 2-IM and show that the integrality gap (i.e., the ratio between the optimal objective value of the ILP and the optimal objective value of its relaxation) is at least $\sqrt{\mu}$, even already for spanning instances with a bipartite union graph. Up to a small constant factor, this equals the (inverse) approximation ratio guaranteed by Algorithm 1, which we will propose in Sect. 3. This serves as a hint that overcoming the approximation dependency $\sqrt{\mu}$ for 2-IM may be hard.

In the context of classical (perfect) matchings, the standard ILP formulation and its LP-relaxation describe the very same feasible points (called *matching polytope*),

which is the corner stone of the problem being solvable in polynomial time [19]. Given a 2-stage graph $\mathcal{G} = (V, E_1, E_2)$, the natural LP-formulation for 2-IM starts with the product of two distinct such perfect matching polytopes. Let $\delta_\ell(v)$ denote all edges incident to vertex v in G_ℓ , and let (M_1, M_2) be a 2-stage perfect matching in \mathcal{G} . For each $\ell \in [2]$, we model M_ℓ via the standard matching polytope: For each $e \in E_\ell$ there is an indicator variable x_e^ℓ that is 1 if and only if $e \in M_\ell$. The constraints (1a) below suffice for bipartite graphs; for general graphs one also considers the *blossom constraints* (1b). These ensure that for any odd-sized vertex set W , at most $\lfloor |W|/2 \rfloor$ edges between those vertices can partake in a matching (see [19] for details). Additionally to these standard descriptions, for each $e \in E_\cap$ we use a variable z_e that is 1 if and only if $e \in M_1 \cap M_2$. We want to maximize $p(M_1, M_2) = \sum_{e \in E_\cap} z_e$, such that:

$$\sum_{e \in \delta_\ell(v)} x_e^\ell = 1 \quad \forall \ell \in [2], \forall v \in V(E_\ell) \tag{1a}$$

$$\sum_{e \in E(G[W])} x_e^\ell \leq \frac{|W| - 1}{2} \quad \forall \ell \in [2], \forall W \subseteq V(E_\ell) \text{ with } |W| \text{ odd} \tag{1b}$$

$$z_e \leq x_e^\ell \quad \forall \ell \in [2], \forall e \in E_\cap \tag{1c}$$

$$x_e^\ell \in \{0, 1\} \quad \forall \ell \in [2], \forall e \in E_\ell \tag{1d}$$

$$z_e \in \{0, 1\} \quad \forall e \in E_\cap \tag{1e}$$

Thereby, constraints (1c), together with the fact that we maximize all z -values, ensure that $z_e = \min\{x_e^1, x_e^2\}$ in any optimal solution.

Theorem 7 *The natural LP for 2-IM has at least an integrality gap of $\sqrt{\mu}$, independent of the number μ of edges in the intersection.*

Proof We construct a family of 2-IM instances, each with bipartite union graph, parameterized by some parameter k . Each instance is reduced and has a maximum profit of 1, but its LP relaxation has objective value at least $k + 1 = \sqrt{\mu}$.

Fix some $k \geq 3$. We construct $\mathcal{G} := \mathcal{G}(k) = (V, E_1, E_2)$ as follows (see Fig. 3 for a visualization with $k = 3$). Let $V := \{a_{i,j}, b_{i,j}\}_{i,j \in [k]} \cup \{c_i, d_i\}_{i \in [k]}$. Let $E_\cap := \{a_{i,j}b_{i,j}\}_{i,j \in [k]}$, i.e., the intersection contains precisely the natural pairings of the a and b vertices. We call these the *shared* edges. In E_1 , we additionally add edges $\{b_{i-1,j}a_{i,j}\}_{i \in [k], j \in [k]}$. Similarly, we add edges $\{b_{i,j-1}a_{i,j}\}_{i \in [k], j \in [k]}$ to E_2 . Now, both stages consist of $k + 1$ disjoint paths of length $2k + 1$ which are “interwoven” between the stages such that (i) every second edge in each path is shared (starting with the first), and (ii) any path in G_1 has exactly one edge in common with every path in G_2 . Let $P_i^\ell, \ell \in [2], i \in [k]$, denote those paths in their natural indexing. We make each stage connected by joining every pair of “neighboring” paths, each together with a c and a d vertex. More precisely, we add edges $\{c_j a_{0,j-1}, c_j a_{0,j}, b_{k,j-1} d_j, b_{k,j} d_j\}_{j \in [k]}$ to E_1 . Analogously, we add $\{c_{\varphi(i)} a_{i-1,0}, c_{\varphi(i)} a_{i,0}, b_{i-1,k} d_{\varphi(i)}, b_{i,k} d_{\varphi(i)}\}_{i \in [k]}$ to E_2 ; the indexing function $\varphi(i) := k - i + 1$ ensures that these new edges are not common to both stages. (In fact, if we would not care for a spanning \mathcal{G} , we could simply use “new” vertices instead of reusing c, d in G_2 .) This finishes the construction, and since \mathcal{G} contains no forbidden edges, it is reduced.




Fig. 3 2-IM instance for $k=3$ with an integrality gap $\geq \sqrt{\mu}$. E_1 is straight and red, E_2 is curvy and blue. Dotted green lines identify vertices (Color figure online)

Since the inner vertices of any path P_i^1 have degree 2 in G_1 , any perfect matching in G_1 either contains all or none of the path’s shared edges. Assume some shared edge $a_{0,j}b_{0,j}$ is in a perfect matching in G_1 . Let C be the path $a_{0,0} c_1 a_{0,1} c_2 \dots c_k a_{0,k}$ in G_1 . Recall that all c -vertices have degree 2 in G_1 . Since $a_{0,j}$ is matched outside of C , all other $a_{0,j'}$, $j' \neq j$, have to be matched with these c -vertices. Thus, P_j^1 is the only path that contributes shared edges to the matching. Conversely, since C contains one less c -vertex than a -vertices, any perfect matching in G_1 has to have at least (and thus exactly) one such path. As the analogous statement holds for G_2 and by the interleaving property (ii) above, any multistage perfect matching contains exactly one shared edge.

However, we construct a feasible fractional solution with objective value $\sqrt{\mu}$.

Let $\lambda := 1/(k + 1)$. We set the x - and z -variables of all shared edges to λ , satisfying all constraints (1c). This uniquely determines all other variable assignments, in order to satisfy (1a): Since the inner vertices of each P_i^ℓ have degree 2 in G_ℓ , the non-shared edges in these path have to be set to $1 - \lambda$. Again consider path C : Each a -vertex in C has an incident shared edge that contributes λ to the sum in the vertex’ constraint (1a); there are no other edges incident to C . Thus, we have to set $x_{c_j a_{0,j-1}}^1 = 1 - j\lambda$ and $x_{c_j a_{0,j}}^1 = j\lambda$ such that, for each vertex in C , its incident variable values sum to 1. The analogous statements holds for the corresponding path through d -vertices in G_1 , and the analogous paths in G_2 . All constraints (1a) are satisfied. The blossom constraints (1b) act only on x -variables, i.e., on individual stages. Since our graph is bipartite, only considering the x -variables of one stage and disregarding (1b) yields the bipartite matching polytope which has only integral vertices; our (sub)solution is an element of this polytope. Thus, (1b) cannot be violated by our assignment.

By construction we have $\mu = (k + 1)^2$. Thus, the objective value of our assignment is $\sum_{e \in E_\cap} \lambda = \mu/(k + 1) = \sqrt{\mu}$, as desired. \square

3 Approximation

We start with the special case of 2-IM, before extending the result to the multistage MIM scenario. Then we will transform the algorithms for use with 2-UM and MUM.

3.1 Approximating 2-IM

We first describe Algorithm 1, which is an approximation for 2-IM. Although its ratio is not constant but grows with the rate of $\sqrt{\mu}$, Theorem 7 hints that better approximations may be hard to obtain. Algorithm 1 roughly works as follows: Given a 2-stage graph \mathcal{G} , we iterate the following procedure on G_1 until every edge of E_\cap has been in at least one perfect matching: Compute a perfect matching M_1 in G_1 that uses the maximum number of edges of E_\cap that have not been used in any previous iteration; then compute a perfect matching M_2 in G_2 that optimizes the profit with respect to M_1 . While doing so, keep track of the maximal occurring profit. Note that by choosing weights appropriately, we can construct a perfect matching that contains the maximum number of edges of some prescribed edge set in polynomial time [19].

We show:

Theorem 8 *Algorithm 1 is a tight $(1/\sqrt{2\mu})$ -approximation for 2-IM.*

We prove this via two Lemmata; the bad instance of Lemma 9 in conjunction with the approximation guarantee (Lemma 10) establishes tightness.

Lemma 9 (Bad instance) *The approximation ratio of Algorithm 1 is at most $1/\sqrt{2\mu}$.*

Proof Consider the following family \mathcal{G}_k of 2-IM instances, parameterized by number $k \geq 1$. An example using $k = 4$ is depicted in Fig. 4. In the first stage, for each $i \in [k]$ create a 4-cycle C_i and label two of its adjacent vertices w'_i and w_i . Add a 3-path with new inner vertices of degree 2 from w_i to w'_{i+1} for each $i \in [k - 1]$. For each $i \in [k - 1]$, create a vertex v_i and an edge $w_i v_i$. Create a vertex u and an edge uw_k . For each $i \in \llbracket k - 1 \rrbracket$, create a vertex u_i and an edge uu_i . For each $i \in [k - 1]$, create a path P_i from u_i to v_i with $2i + 1$ edges and label the new inner vertices with $a_1^i, b_1^i, a_2^i, b_2^i, \dots, a_i^i, b_i^i$ in this order.

The second stage is constructed isomorphically to the first stage. To avoid ambiguity in the naming, we underline element names of the second stage. The 2-stage graph \mathcal{G}_k is

Algorithm 1: Approximation of 2-IM

```

1 set  $(M_1, M_2)$  to  $(\emptyset, \emptyset)$ 
2 for  $i = 1, 2, \dots$  do
3   define weights  $w_1 : E_1 \rightarrow \{0, 1\}$  to be 1 exactly for  $e \in E_\cap \setminus \bigcup_{j \in [i-1]} M_1^{(j)}$ 
4   compute a maximum weight perfect matching  $M_1^{(i)}$  in  $G_1$  w.r.t.  $w_1$ 
5   define weights  $w_2 : E_2 \rightarrow \{0, 1\}$  to be 1 exactly for  $e \in M_1^{(i)}$ 
6   compute a maximum weight perfect matching  $M_2^{(i)}$  in  $G_2$  w.r.t.  $w_2$ 
7   if  $|M_1^{(i)} \cap M_2^{(i)}| \geq |M_1 \cap M_2|$  then set  $(M_1, M_2)$  to  $(M_1^{(i)}, M_2^{(i)})$  if  $E_\cap \subseteq \bigcup_{j \in [i]} M_1^{(j)}$  then
   return  $(M_1, M_2)$ 

```




Fig. 4 2-IM instance \mathcal{G}_4 as in Lemma 9. Edges in E_1 are curvy and blue, edges in E_2 straight and red. The vertices are labeled according to the first stage (Color figure online)

completely defined by the following identifications: for each $i \in [k]$, let $w'_i = w_{k-i+1}$ and $w_i = w'_{k-i+1}$; for each $i \in [k - 1]$ and each $j \in [i]$, let $a'_j = b_{k-i}^{k-j}$ and $b'_j = a_{k-i}^{k-j}$. Thus, $E_\cap = F \cup A$ is precisely the union of $F := \{w_i w'_i \mid i \in [k]\}$ and $A := \{a'_j b'_j \mid i \in [k - 1], j \in [i]\}$. Observe that \mathcal{G}_k is reduced, its union graph is bipartite, and $\mu = k + \sum_{i \in [k-1]} i = k(k + 1)/2$.

By construction, for any perfect matching M in the first (second) stage, $|M \cap E_\cap| \leq k$. Let $M_F (\underline{M}_F)$ denote the unique perfect matching that contains uu_0 ($u\underline{u}_0$, respectively) and all of F . The pair (M_F, \underline{M}_F) is an optimal solution with profit $|F| = k$.

Consider an alternative perfect matching M in the first stage. In each cycle C_i , we consider the shared edge $w_i w'_i$ and its *opposing edge* (i.e., its unique non-adjacent edge in C_i). We distinguish between three possibilities regarding their memberships in M : the shared edge and its opposing edge are in M (*type Y*), only the opposing edge is in M (*type N1*), none of them are in M (*type N2*).

Picking an edge adjacent to u determines a perfect matching up to the types of some C_i -cycles. For $i \in \llbracket k - 1 \rrbracket$, let M_i denote the unique perfect matching that contains uu_i , is type Y in C_{i+1} , but type $N2$ in each C_j for $j > i + 1$. Note that M_i is type $N1$ in each C_j with $j \leq i$ and contains all i shared edges along P_i . Thus, $|M_i \cap E_\cap| = i + 1$. Most importantly, consider any perfect matching M' in the second stage. By construction, for any $i \in \llbracket k - 1 \rrbracket$, no two edges of $(M_i \cap A) \cup \{w'_{i+1} w_{i+1}\}$ can be contained simultaneously in M' . It follows that $|M_i \cap M'| \leq 1$,

Algorithm 1 may never choose the optimal M_F as a perfect matching for the first stage: In the first iteration, both M_F and M_{k-1} have weight k , so the algorithm may choose M_{k-1} and obtain a 2-stage perfect matching with profit 1. In the following iteration, the weight (denoting the preference of edges) of M_F is decreased by 1, since the edge $w'_k w_k$ has already been chosen in M_{k-1} . Consequently, in each following iteration $i \in [k]$ the algorithm may choose M_{k-i} over M_F , each time decreasing the weight of M_F by 1. After choosing M_0 over M_F in iteration k , each edge in E_\cap has




Fig. 5 Visualization of the relationships between $E_\cap, M_\cap^*, M_1^{(i)}, R_i$ and R_i^* for $i \in [3]$

been in some matching in the first stage; the algorithm stops and returns a 2-stage perfect matching with profit 1.

Since $\mu = k(k + 1)/2$, the optimal profit is $k = (\sqrt{8\mu + 1} - 1)/2$. Thus, the approximation factor is at most $1/k = 2/(\sqrt{8\mu + 1} - 1)$ which tends to our guarantee of $1/\sqrt{2\mu}$ for increasing k . \square

Lemma 10 (Guarantee) *The approximation ratio of Algorithm 1 is at least $1/\sqrt{2\mu}$.*

Proof Let \mathcal{G} be a feasible and reduced 2-stage graph with non-empty E_\cap . Clearly, our algorithm achieves $\text{apx} \geq 1$ as described in Observation 3. Let k denote the number of iterations. For any $i \in [k]$, let $(M_1^{(i)}, M_2^{(i)})$ denote the 2-stage perfect matching computed in the i th iteration. The algorithm picks at least one new edge of E_\cap per iteration into $M_1^{(i)}$ and hence terminates. Let (M_1^*, M_2^*) denote an optimal 2-stage perfect matching and $M_\cap^* := M_1^* \cap M_2^*$ its intersection. Let $R_i := (M_1^{(i)} \cap E_\cap) \setminus \bigcup_{j \in [i-1]} R_j$ denote the set of edges in $M_1^{(i)} \cap E_\cap$ that are not contained in $M_1^{(j)}$ for any previous iteration $j < i$ and let $r_i := |R_i|$. Note that in iteration i , the algorithm first searches for a perfect matching $M_1^{(i)}$ in G_1 that maximizes the cardinality r_i of its intersection with $E_\cap \setminus \bigcup_{j \in [i-1]} R_j$. We define $R_i^* := (M_1^{(i)} \cap M_\cap^*) \setminus \bigcup_{j \in [i-1]} R_j^*$ and $r_i^* := |R_i^*|$ equivalently to R_i , but w.r.t. M_\cap^* (see Fig. 5). Observe that $R_i \cap M_\cap^* = R_i^*$.

Let $q := \sqrt{2\mu}$. For every $i \in [k]$ the algorithm chooses $M_2^{(i)}$ such that $|M_1^{(i)} \cap M_2^{(i)}|$ is maximized. Since we may choose $M_2^{(i)} = M_2^*$, it follows that $\text{apx} \geq \max_{i \in [k]} r_i^*$. Thus, if $\max_{i \in [k]} r_i^* \geq \text{opt}/q$, we have a $(1/q)$ -approximation. In case $\text{opt} \leq q$, any solution with profit at least 1 yields a $(1/q)$ -approximation. We show that we are in one of these cases.

Let $\bar{q} := \lceil q \rceil$. Assume that $\text{opt} > q$ (thus $\text{opt} \geq \bar{q}$) and simultaneously $r_i^* < \text{opt}/q$ for all $i \in [k]$. Since we distribute M_\cap^* over the disjoint sets $\{R_i^* \mid i \in [k]\}$, each containing less than opt/q edges, we know that $k > q$ (thus $k \geq \bar{q}$). In iteration i , M_1^* has weight $|(M_1^* \cap E_\cap) \setminus \bigcup_{j \in [i-1]} R_j| \geq |M_\cap^* \setminus \bigcup_{j \in [i-1]} R_j| = |M_\cap^* \setminus \bigcup_{j \in [i-1]} R_j^*|$. Hence, the latter term is a lower bound on r_i , that we estimate as follows: $r_i \geq |M_\cap^* \setminus \bigcup_{j \in [i-1]} R_j^*| = \text{opt} - \sum_{j \in [i-1]} r_j^* \geq \text{opt} - \sum_{j \in [i-1]} \text{opt}/q = \text{opt} \cdot (1 - (i - 1)/q)$. The above assumptions give a contradiction:

Algorithm 2: General multistage approximation

Input: Temporal graph \mathcal{G} , 2-stage perfect matching algorithm \mathcal{A}

- 1 create path $P := \{e_1, \dots, e_{\tau-1}\}$
- 2 **foreach** $i \in [\tau - 1]$ **do**
- 3 set (S_i, T_{i+1}) to $\mathcal{A}(V, E_i, E_{i+1})$ // approximate 2-stage graphs
- 4 set weight of e_i to $w_i := |S_i \cap T_{i+1}|$
- 5 compute maximum weight matching M_P in (P, w)
- 6 set $(M_i)_{i \in [\tau-1]}$ to $(S_i)_{i \in [\tau-1]}$ and M_τ to T_τ // set initial solution
- 7 **foreach** $i \in [\tau - 1]$ **do** // modify solution according to M_P
- 8 **if** $e_i \in M_P$ **then** set M_i to S_i and M_{i+1} to T_{i+1}
- 9 **return** (M_1, \dots, M_τ)

$$\begin{aligned} \mu &= \left| \bigcup_{i \in [k]} R_i \right| \geq \sum_{i \in [\bar{q}]} r_i \geq \text{opt} \cdot \sum_{i \in [\bar{q}]} \left(1 - \frac{i-1}{q}\right) \geq \bar{q} \cdot \sum_{i \in [\bar{q}]} \left(1 - \frac{i-1}{q}\right) \\ &= \bar{q} \left(\bar{q} - \sum_{i \in [\bar{q}-1]} \frac{i}{q}\right) \\ &= \bar{q} \left(\bar{q} - \frac{(\bar{q}-1)\bar{q}}{2q}\right) = \bar{q}^2 \left(1 - \frac{\bar{q}-1}{2q}\right) > \bar{q}^2 \left(1 - \frac{q}{2q}\right) \geq \mu. \end{aligned}$$

□

3.2 Approximating MIM

Let us extend the above result to an arbitrary number of stages. We show that we can use any 2-IM approximation algorithm (in particular also Algorithm 1) as a black box to obtain an approximation algorithm for MIM, while only halving the approximation ratio: Algorithm 2 uses an edge weighted path (P, w) on τ vertices as an auxiliary graph. We set the weight of the edge between the i th and $(i + 1)$ th vertex to an approximate solution for the 2-IM instance that arises from the i th and $(i + 1)$ th stage of the MIM instance. A maximum weight matching M_P in (P, w) induces a feasible solution for the MIM problem: If an edge $(j, j + 1)$ is in M_P , we use the corresponding solutions for the j th and $(j + 1)$ th stage; for stages without incident edge in M_P , we select an arbitrary solution. Since no vertex is incident to more than one edge in M_P , there are no conflicts.

Observation 11 For $F \subseteq E(P)$, denote $w(F) := \sum_{e \in F} w(e)$. Let e_i denote the i th edge of P . For $b \in [2]$, the matchings $M_b := \{e_i \in E(P) \mid i = b \pmod 2\}$ are disjoint and their union is exactly $E(P)$. Thus, any maximum weight matching M_P in P achieves $2 \cdot w(M_P) \geq w(E(P))$.

Theorem 12 For a 2-IM α -approximation, Algorithm 2 $\frac{\alpha}{2}$ -approximates MIM.

Proof Let $\mathcal{G} = (V, E_1, \dots, E_\tau)$ be the given temporal graph. For any $i \in [\tau - 1]$, (S_i, T_{i+1}) is the output of the 2-IM α -approximation $\mathcal{A}(V, E_i, E_{i+1})$; let $w_i := |S_i \cap T_{i+1}|$. Let $\mathcal{M}^* := (M_1^*, \dots, M_\tau^*)$ denote a multistage perfect matching whose

profit $p(\mathcal{M}^*)$ is maximum. Since \mathcal{A} is an α -approximation for 2-IM, we know that $|M_i^* \cap M_{i+1}^*| \leq w_i/\alpha$ for every $i \in [\tau - 1]$. Thus $p(\mathcal{M}^*) \leq (1/\alpha) \sum_{i \in [\tau-1]} w_i$. Algorithm 2 computes a maximum weight matching M_P in (P, w) and constructs a multistage solution \mathcal{M} . By Observation 11, we obtain $p(\mathcal{M}^*) \leq (1/\alpha) \sum_{i \in [\tau-1]} w_i = (1/\alpha) \cdot w(E(P)) \leq (2/\alpha) \cdot w(M_P) \leq (2/\alpha) \cdot p(\mathcal{M})$. \square

We compute a maximum weight matching in a path in linear time using straightforward dynamic programming. Hence, assuming running time f for \mathcal{A} , Algorithm 2 requires $\mathcal{O}(\sum_{i \in [\tau-1]} |f(G_i, G_{i+1})|)$ steps.

Corollary 13 *Algorithms 1 in Algorithm 2 yields a $(1/\sqrt{8\mu})$ -approximation for MIM.*

There is another way to approximate MIM via an approximation for 2-IM, which neither dominates nor is dominated by the above method:

Theorem 14 *There is an S-reduction from MIM to 2-IM, i.e., given any MIM instance \mathcal{G} , we can find a corresponding 2-IM instance \mathcal{G}' in polynomial time such that any solution for \mathcal{G} bijectively corresponds to a solution for \mathcal{G}' with the same profit. Furthermore, $|E(G'_1) \cap E(G'_2)| = \sum_{i \in [\tau-1]} |E(G_i) \cap E(G_{i+1})|$.*

Proof We will construct a 2-stage graph \mathcal{G}' whose first stage G'_1 consists of (subdivided) disjoint copies of G_i for odd i ; conversely its second stage G'_2 consists of (subdivided) disjoint copies of G_i for even i . More precisely, consider the following construction: Let $b(i) := 2 - (i \bmod 2)$. For each $i \in [\tau]$, we create a copy of G_i in $G'_{b(i)}$ where each edge $e \in E(G_i)$ is replaced by a 7-path p_i^e . We label the 3rd (5th) edge along p_i^e (disregarding its orientation) with e_i^- (e_i^+ , respectively). To finally obtain \mathcal{G}' , for each $i \in [\tau - 1]$ and $e \in E(G_i) \cap E(G_{i+1})$, we identify the vertices of e_i^+ with those of e_{i+1}^- (disregarding the edges' orientations); thereby precisely the edges e_i^+ and e_{i+1}^- become an edge common to both stages. No other edges are shared between both stages. This completes the construction of \mathcal{G}' and we have $|E(G'_1) \cap E(G'_2)| = \sum_{i \in [\tau-1]} |E(G_i) \cap E(G_{i+1})|$.

Assume $\mathcal{M}' := (M'_1, M'_2)$ is a solution for \mathcal{G}' . Clearly, each path p_i^e in $G'_{b(i)}$ is matched alternately and hence either all or none of e_i^-, e_i^+ , the first, and the last edge of p_i^e are in $M'_{b(i)}$. We derive a corresponding solution \mathcal{M} for \mathcal{G} : For every $i \in [\tau]$ and $e \in E(G_i)$, we add e to M_i if and only if $e_i^- \in M'_{b(i)}$. Suppose that M_i is not a perfect matching for G_i , i.e., there exists a vertex v in G_i that is not incident to exactly one edge in M_i . Then also the copy of v in the copy of G_i in $G'_{b(i)}$ is not incident to exactly one edge of $M'_{b(i)}$, contradicting the feasibility of \mathcal{M}' .

Consider the profit achieved by \mathcal{M} : Every edge in $M'_1 \cap M'_2$ corresponds to a different identification $\langle e_i^+, e_{i+1}^- \rangle$. We have $e \in M_i \cap M_{i+1}$ if and only if $e_i^- \in M'_{b(i)}$, $e_{i+1}^- \in M'_{b(i+1)}$, and $e_i^+ = e_{i+1}^-$. It follows that this holds if and only if $e_i^+ \in M'_{b(i)} \cap M'_{b(i+1)}$ and hence the profit of \mathcal{M} is equal to that of \mathcal{M}' . The inverse direction proceeds in the same manner. \square

Since the new $\mu' := |E(G'_1) \cap E(G'_2)|$ is largest w.r.t. the original μ if $|E(G_i) \cap E(G_{i+1})|$ is constant for all i , we obtain:

Corollary 15 *For any 2-IM $\alpha(\mu)$ -approximation where $\alpha(\mu)$ is a (typically decreasing) function of μ , there is an $\alpha((\tau - 1)\mu)$ -approximation for MIM. Using Algorithm 1, this yields a ratio of $1/\sqrt{2(\tau - 1)\mu}$; for 3-IM and 4-IM this is tighter than Theorem 12.*

Assume the approximation ratio for 2-IM would not depend on μ . Then the above would yield a surprisingly strong result:

Corollary 16 *Any 2-IM α -approximation with constant α results in an α -approximation of MIM. If MIM is APX-hard, so is 2-IM.*

3.3 Approximating MUM

Consider the MUM-problem which minimizes the cost. As noted in Observation 4, a 2-approximation is easily accomplished. However, by exploiting the previous results for MIM, we obtain better approximations.

Theorem 17 *Any α -approximation of MIM is a $(2 - \alpha)$ -approximation of MUM.*

Proof Recall that an optimal solution of MIM constitutes an optimal solution of MUM. As before, we denote the heuristic sequence of matchings by $(M_i)_{i \in [\tau]}$ and the optimal one by $(M_i^*)_{i \in [\tau]}$. Let $\xi := \sum_{i \in [\tau-1]} (n_i + n_{i+1})/2$. Consider the solutions' values w.r.t. MUM:

$$\begin{aligned} \frac{\text{apx}_U}{\text{opt}_U} &= \frac{\sum_{i \in [\tau-1]} c(M_i, M_{i+1})}{\sum_{i \in [\tau-1]} c(M_i^*, M_{i+1}^*)} \\ &= \frac{\xi - \sum_{i \in [\tau-1]} |M_i \cap M_{i+1}|}{\xi - \sum_{i \in [\tau-1]} |M_i^* \cap M_{i+1}^*|} \leq \frac{\xi - \alpha \cdot \text{opt}_\cap}{\xi - \text{opt}_\cap} =: f. \end{aligned}$$

As $0 < \alpha < 1$, f is monotonously increasing in opt_\cap if $0 \leq \text{opt}_\cap < \xi$. Thus, since $\text{opt}_\cap \leq \sum_{i \in [\tau-1]} \min(n_i, n_{i+1})/2 \leq \sum_{i \in [\tau-1]} (n_i + n_{i+1})/4 = \xi/2$, it follows that $\text{apx}_U/\text{opt}_U \leq (\xi - \alpha \cdot \xi/2)/(\xi - \xi/2) = 2 - \alpha$. \square

Corollary 18 *Let $r := \min\{8, 2(\tau - 1)\}$. We have a $(2 - 1/\sqrt{r \cdot \mu})$ -approximation for MUM.*

Note that a similar reduction from MIM to MUM is not achieved as easily: Consider any $(1 + \varepsilon)$ -approximation for MUM. Choose an even integer $k \geq 6$ such that $k/(k - 1) \leq 1 + \varepsilon$; consider a spanning 2-stage instance where each stage is a k -cycle and E_\cap consists of a single edge e . The optimal 2-stage perfect matching \mathcal{M}^* that contains e in both stages has profit $p(\mathcal{M}^*) = 1$ and cost $c(\mathcal{M}^*) = 2 \cdot k/2 - 1 = k - 1$. A 2-stage perfect matching \mathcal{M} that does not contain e still satisfies $c(\mathcal{M}) = k$ and as such is an $(1 + \varepsilon)$ -approximation for MUM. However, its profit $p(\mathcal{M}) = 0$ does not provide any approximation of $p(\mathcal{M}^*) = 1$.

As for MIM, we aim to extend a given approximation for 2-UM to a general approximation for MUM. Unfortunately, we cannot use Theorems 14 and 17 for this, as an approximation for 2-UM does not generally constitute one for 2-IM (and MIM). On the positive side, a similar approach as used in the proof of Theorem 12 also works for minimization.

Theorem 19 Any α -approximation \mathcal{A} for 2-UM results in a $(1 + \alpha/2)$ -approximation for MUM by using \mathcal{A} in Algorithm 2.

Proof As before, let $(M_i^*)_{i \in [\tau]}$ denote an optimal solution for MUM. For each $i \in [\tau - 1]$, (S_i, T_i) denotes the output of $\mathcal{A}(V, E_i, E_{i+1})$. For $L \subseteq [\tau - 1]$, let $\xi(L) := \sum_{i \in L} (n_i + n_{i+1})/2$ and $\sigma(L) := \sum_{i \in L} |S_i \cup T_i|$. Note that $w_i := \xi(\{i\}) - \sigma(\{i\})$ equals the weight of e_i . We define $I := \{i \in [\tau - 1] \mid e_i \in M_P\}$ as the set of indices corresponding to M_P and $J := [\tau - 1] \setminus I$ as its complement. By Observation 11, we have $w(E(P)) \leq 2 \cdot w(M_P)$, thus

$$\begin{aligned} \xi(I) - \sigma(I) + \xi(J) - \sigma(J) &= w(E(P)) \leq 2 \cdot w(M_P) = 2(\xi(I) - \sigma(I)) \\ \Rightarrow \sigma(I) + \xi(J) &\leq \xi(I) + \sigma(J) \Rightarrow 2(\sigma(I) + \xi(J)) \leq \xi(I \cup J) + \sigma(I \cup J). \end{aligned}$$

The trivial upper bound ξ suffices to bound the algorithm's solution value:

$$\text{apx} = \sigma(I) + \sum_{j \in J} |M_j \cup M_{j+1}| \leq \sigma(I) + \xi(J) \leq \frac{1}{2}(\xi(I \cup J) + \sigma(I \cup J)).$$

Since $\sigma(I \cup J)$ α -approximates the sum of all 2-UM instances' solution values, we have $\sigma(I \cup J) \leq \alpha \cdot \text{opt}$. For each transition, any solution satisfies $(n_i + n_{i+1})/4 \leq |M_i \cup M_{i+1}|$ and hence $\xi(I \cup J) \leq 2 \cdot \text{opt}$. Finally, we obtain the claimed ratio: $\text{apx} \leq 1/2 \cdot (2 \cdot \text{opt} + \alpha \cdot \text{opt}) = (1 + \alpha/2) \cdot \text{opt}$. \square

4 Conclusion

In this paper we presented the first approximation algorithm for 2-IM with a (tightly analyzed) approximation ratio of $1/\sqrt{2\mu}$. It remains open if a constant factor approximation for 2-IM is possible; however, we showed that this would imply a constant factor approximation for MIM. We further showed two ways in which MIM and MUM can be approximated by using any algorithm that approximates 2-IM, thereby also presenting the first approximation algorithms for multistage matching problems with an arbitrary number of stages.

Furthermore, it turns out that the techniques of Algorithms 1 and 2 for MIM can be used as a building block to approximate a broad set of related problems, so-called *Multistage Subgraph Problems* [12].

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix

A A Note on Approximating Max-MPM

The proposed $(1/2)$ -approximation for Max-MPM in [2] does not work. It takes a temporal graph as input, where each stage may be an arbitrary graph (not necessarily complete), picks a matching for every second stage G_i , and reuses the same matching for stage G_{i+1} . Thus, every second stage transition is optimal, whereas every other second transition potentially constitutes a worst case. If the algorithm’s solution is feasible, we indeed yield the proposed approximation ratio. However, such an approach is inherently problematic as there is no reason why a matching in G_i would need to be feasible for G_{i+1} . In fact, consider a temporal graph $\mathcal{G} = (V, E_1, \dots, E_\tau)$ with $V = \{v_1, \dots, v_4\}$. Let $E_i = \{v_1v_2, v_3v_4\}$ for odd i , and $E_i = \{v_2v_3, v_4v_1\}$ for even i . No perfect matching in E_i is also a perfect matching in E_{i+1} .

Thus, although any α -approximation for Max-MPM would directly yield an α -approximation for MIM on spanning temporal graphs, we currently do not know of any such algorithm. In fact, a constant-factor approximation seems difficult to obtain, see Theorem 7. Personal communication with B. Escoffier confirmed our counterexample. One may consider a relaxed version of Max-MPM where one tries to find matchings of large weight in each stage, formally optimizing the weighted sum between the profit and the summed stagewise matching weights. Observe that in this scenario it is not guaranteed that the optimal solution induces a perfect (nor even maximum) matching in each stage. However, for this problem their analysis would be correct and their algorithm yields a $(1/2)$ -approximation.

Counterexample

We examine the $(1/2)$ -approximation algorithm \mathcal{A} for Max-MPM that was proposed in [2, Theorem 8], and give a reduced spanning 4-stage instance where \mathcal{A} does not yield a feasible solution. We use four stages since the algorithm treats fewer stages as special cases. Still, the feasibility problem that we are about to describe is inherent to all its variants.

Consider the temporal graph $\mathcal{G} = (V, E_1, E_2, E_3, E_4)$ given in Fig. 6, where $E_1 = E_3$ and $E_2 = E_4$ and edges have uniform weight 0. We trivially observe that any perfect matching in G_i is optimal w.r.t. edge weight. For each $i \in [3]$, we have $E_i \cap E_{i+1} = \{e_1, e_2, e_3, e_4\}$.

The algorithm proceeds as follows on \mathcal{G} : For each $i \in [3]$, it computes a perfect matching M_i in G_i that maximizes $|M_i \cap E_{i+1}|$. It constructs the solutions $\mathcal{M} := (M_1, M_1, M_3, M_3)$ and $\mathcal{M}' := (\hat{M}_1, M_2, M_2, \hat{M}_3)$, where \hat{M}_i is an arbitrary perfect matching in G_i , and outputs the solution that maximizes the profit.

Any perfect matching M_1 in G_1 that maximizes $|M_1 \cap E_2|$ must contain both e_1 and e_2 and as such also f_1 . This contradicts the feasibility of \mathcal{M} , since $f_1 \notin E_2$. Conversely, any such perfect matching M_2 in G_2 must contain both e_3 and e_4 and as such also f_2 . Again, this contradicts the feasibility of \mathcal{M}' , since $f_2 \notin E_3$. It follows that the algorithm cannot pick a feasible solution.




Fig. 6 Counterexample for the proposed $(1/2)$ -approximation. Edges in $E_1=E_3$ are curly (and blue), edges in $E_2=E_4$ are straight (and red) (Color figure online)

We are not aware of any way to circumvent this problem.

References

1. Akrida, E.C., Mertzios, G.B., Spirakis, P.G., Zamarayev, V.: Temporal vertex cover with a sliding time window. *J. Comput. Syst. Sci.* **107**, 108–123 (2020). <https://doi.org/10.1016/j.jcss.2019.08.002>
2. Bampis, E., Escoffier, B., Lampis, M., Paschos, V.T.: Multistage matchings. In: 16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018), vol. 101, pp. 7:1–7:13 (2018). <https://doi.org/10.4230/LIPIcs.SWAT.2018.7>
3. Bampis, E., Escoffier, B., Schewior, K., Teiller, A.: Online multistage subset maximization problems. In: 27th Annual European Symposium on Algorithms (ESA 2019), vol. 144, pp. 11:1–11:14 (2019). <https://doi.org/10.4230/LIPIcs.ESA.2019.11>
4. Bampis, E., Escoffier, B., Teiller, A.: Multistage knapsack. In: 44th International Symposium on Mathematical Foundations of Computer Science (MFCS 2019), vol. 138, pp. 22:1–22:14 (2019). <https://doi.org/10.4230/LIPIcs.MFCS.2019.22>
5. Bampis, E., Escoffier, B., Kononov, A.V.: Lp-based algorithms for multistage minimization problems. In: Approximation and Online Algorithms: 18th International Workshop, WAOA 2020, Virtual Event, September 9–10, 2020, Revised Selected Papers, pp. 1–15 (2020). https://doi.org/10.1007/978-3-030-80879-2_1
6. Baste, J., Bui-Xuan, B.M., Roux, A.: Temporal matching. *Theor. Comput. Sci.* **806**, 184–196 (2020). <https://doi.org/10.1016/j.tcs.2019.03.026>
7. Bernstein, A., Stein, C.: Fully dynamic matching in bipartite graphs. In: Proceedings of 42nd International Colloquium on Automata, Languages and Programming (ICALP 2015), pp. 167–179 (2015). https://doi.org/10.1007/978-3-662-47672-7_14
8. Bhattacharya, S., Henzinger, M., Italiano, G.F.: Deterministic fully dynamic data structures for vertex cover and matching. *SIAM J. Comput.* **47**(3), 859–887 (2018). <https://doi.org/10.1137/140998925>
9. Bosek, B., Leniowski, D., Sankowski, P., Zych, A.: Online bipartite matching in offline time. In: 2014 IEEE 55th Annual Symposium on Foundations of Computer Science, pp. 384–393 (2014). <https://doi.org/10.1109/FOCS.2014.48>
10. Bredereck, R., Komusiewicz, C., Kratsch, S., Molter, H., Niedermeier, R., Sorge, M.: Assessing the computational complexity of multilayer subgraph detection. *Netw. Sci.* **7**(2), 215–241 (2019). <https://doi.org/10.1017/nws.2019.13>
11. Casteigts, A.: A journey through dynamic networks (with excursions). Habilitation, Université de Bordeaux (2018)
12. Chimani, M., Troost, N., Wiedera, T.: A general approach to approximate multistage subgraph problems. *CoRR* (2021). [arXiv:2107.02581](https://arxiv.org/abs/2107.02581)
13. Eppstein, D.: Offline algorithms for dynamic minimum spanning tree problems. In: Algorithms and Data Structures (WADS 1991), pp. 392–399 (1991). <https://doi.org/10.1007/BFb0028278>
14. Fluschnik, T., Niedermeier, R., Rohm, V., Zschoche, P.: Multistage vertex cover. In: 14th International Symposium on Parameterized and Exact Computation (IPEC 2019), vol. 148, pp. 14:1–14:14 (2019). <https://doi.org/10.4230/LIPIcs.IPEC.2019.14>
15. Garey, M.R., Johnson, D.: Computers and Intractability: A Guide to the Theory of NP-Completeness, W. H Freeman (1979)

16. Gupta, A., Talwar, K., Wieder, U.: Changing bases: multistage optimization for matroids and matchings. In: Proceedings of 41st International Colloquium on Automata, Languages and Programming (ICALP 2014) (2014). https://doi.org/10.1007/978-3-662-43948-7_47
17. Heeger, K., Himmel, A.S., Kammer, F., Niedermeier, R., Renken, M., Sajenko, A.: Multistage problems on a global budget. *Theor. Comput. Sci.* **868**, 46–64 (2021). <https://doi.org/10.1016/j.tcs.2021.04.002>
18. Kempe, D., Kleinberg, J.M., Kumar, A.: Connectivity and inference problems for temporal networks. In: Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing (STOC 2000), pp. 504–513 (2000). <https://doi.org/10.1145/335305.335364>
19. Lovász, L., Plummer, M.: *Matching Theory*, American Mathematical Society (1986)
20. Mertzos, G.B., Molter, H., Niedermeier, R., Zamaraev, V., Zschoche, P.: Computing maximum matchings in temporal graphs. In: 37th International Symposium on Theoretical Aspects of Computer Science (STACS 2020), pp. 27:1–27:14 (2020). <https://doi.org/10.4230/LIPIcs.STACS.2020.27>
21. Michail, O., Spirakis, P.G.: Traveling salesman problems in temporal graphs. In: *Mathematical Foundations of Computer Science 2014 (MFCS 2014)* (2014). https://doi.org/10.1007/978-3-662-44465-8_47
22. Rabin, M.O., Vazirani, V.V.: Maximum matchings in general graphs through randomization. *J. Algorithms* (1989). [https://doi.org/10.1016/0196-6774\(89\)90005-9](https://doi.org/10.1016/0196-6774(89)90005-9)
23. Sankowski, P.: Faster dynamic matchings and vertex connectivity. In: *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2007)*, pp. 118–126 (2007)
24. Thorup, M.: Near-optimal fully-dynamic graph connectivity. In: *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing (STOC 2000)*, pp. 343–350 (2000). <https://doi.org/10.1145/335305.335345>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.