Stable Fractional Matchings

IOANNIS CARAGIANNIS, University of Patras & CTI "Diophantus", Greece
ARIS FILOS-RATSIKAS, École polytechnique fédérale de Lausanne, Switzerland
PANAGIOTIS KANELLOPOULOS, University of Patras & CTI "Diophantus", Greece
ROHIT VAISH, Rensselaer Polytechnic Institute, USA

We study a generalization of the classical stable matching problem that allows for *cardinal* preferences (as opposed to ordinal) and *fractional* matchings (as opposed to integral). After observing that, in this cardinal setting, stable fractional matchings can have much higher social welfare than stable integral ones, our goal is to understand the computational complexity of finding an optimal (i.e., welfare-maximizing) or nearly-optimal stable fractional matching. We present simple approximation algorithms for this problem with weak welfare guarantees and, rather unexpectedly, we furthermore show that achieving better approximations is hard. This computational hardness persists even for approximate stability. To the best of our knowledge, these are the first computational complexity results for stable fractional matchings. En route to these results, we provide a number of structural observations.

CCS Concepts: • **Theory of computation → Design and analysis of algorithms**; **Algorithmic game theory**; *Approximation algorithms analysis.*

Additional Key Words and Phrases: Stable Matchings; Cardinal Preferences; Welfare Maximization.

ACM Reference Format:
Ioannis Caragiannis, Aris Filos-Ratsikas, Panagiotis Kanellopoulos, and Rohit Vaish. 2019. Stable Fractional Matchings. In *ACM EC '19: ACM Conference on Economics and Computation (EC '19), June 24–28, 2019, Phoenix, AZ, USA.* ACM, New York, NY, USA, 19 pages. https://doi.org/10.1145/3328526.3329637

1 INTRODUCTION

The stable matching problem [16] is one of the most extensively studied problems at the interface of economics and computer science, with notable practical applications such as matching students to schools [1], medical graduates to hospitals [31], and organ donors to patients [33]. The input to the problem consists of the preference lists of two sets of agents, commonly referred to as the *men* and the *women*. The goal is to find a *stable* matching, i.e., a matching in which no pair of man and woman prefer each other over their assigned partners.

The standard formulation of the stable matching problem involves two key assumptions: First, that the matching is *integral* (i.e., two agents are either completely matched or completely unmatched), and second, that agents have *ordinal* preferences (typically in the form of rank-ordered lists). Although these assumptions suffice in a number of applications (including those mentioned above), there are natural examples where they turn out to be inadequate. For instance, time-sharing applications [32] naturally give rise to fractional matchings: Imagine a scenario where a set of newly hired employees are matched with a set of supervisors. Assuming that each individual can spend one unit of time at work, an *integral* matching prescribes that every employee should work

EC '19, June 24–28, 2019, Phoenix, AZ, USA

ACM ISBN 978-1-4503-6792-9/19/06…$15.00
https://doi.org/10.1145/3328526.3329637

full-time with a single supervisor. On the other hand, *fractional* matchings allow the employees to divide their time in working with multiple supervisors, making them a more natural modeling choice in such situations.

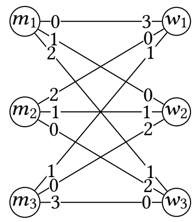

Fig. 1. *Instance with cardinal preferences.*

Likewise, *ordinal* preferences, despite their simplicity and ease of elicitation, can often be quite restrictive. Indeed, in many real-world matching applications (e.g., labor markets), the outcomes experienced by the participants are inherently *cardinal* in nature (e.g., wages). In such settings, it is decidedly more natural to model the *intensity* of preferences, as has been noted in theory [3, 30] as well as in lab experiments [13]. Encouraged by these insights, and the fact that cardinal utilities provide a clean and unambiguous way of comparing matching outcomes in terms of their social welfare, we consider a generalization of the stable matching model that allows for *fractional* matchings (as opposed to integral) and *cardinal* preferences (as opposed to ordinal).

More concretely, we consider a setting where the preferences are specified in terms of *valuations* (e.g., in the matching instance in Figure 1, m_1 values w_1 at 0 and w_1 values m_1 at 3). A fractional matching is simply a convex combination of integral matchings, and an agent's utility under a fractional matching is the appropriately weighted sum of its utilities under the constituent integral matchings. A fractional matching μ is *stable* if no pair of man and woman simultaneously derive greater utility in being integrally matched to each other than they do under μ.

The above generalization has a clear merit in terms of social welfare: Stable solutions under the generalized model can have greater welfare than those under the standard model, as the following example illustrates.

Example 1. Consider the instance in Figure 1 with three men m_1, m_2, m_3 and three women w_1, w_2, w_3. Among the six possible integral matchings, only $\mu_1 := \{(m_1, w_3), (m_2, w_2), (m_3, w_1)\}$ and $\mu_2 := \{(m_1, w_1), (m_2, w_3), (m_3, w_2)\}$ are stable. Indeed, μ_1 and μ_2 are the men-proposing and women-proposing Gale-Shapley matchings respectively [16]. The social welfare (i.e., sum of utilities of all agents) under these matchings is $\mathcal{W}(\mu_1) = \mathcal{W}(\mu_2) = 7$.

Define $\mu_3 := \{(m_1, w_1), (m_2, w_2), (m_3, w_3)\}$, and notice that $\mathcal{W}(\mu_3) = 8$. Now consider a *fractional* matching $\mu := \frac{1}{2}\mu_2 + \frac{1}{2}\mu_3$. The social welfare of μ is $\mathcal{W}(\mu) = \frac{1}{2}\mathcal{W}(\mu_2) + \frac{1}{2}\mathcal{W}(\mu_3) = 15/2 > 7 = \mathcal{W}(\mu_1) = \mathcal{W}(\mu_2)$, i.e., μ has a higher social welfare than any stable integral matching. Importantly, μ is a *stable fractional* matching. Indeed, under μ, the utilities of m_1, m_2, m_3, w_1, w_2, and w_3 are 0, $\frac{1}{2}$, $\frac{3}{2}$, 3, $\frac{3}{2}$, and 1 respectively. Thus, for every man-woman pair, at least one of the two agents meets the corresponding utility threshold for that pair, implying that μ is stable.

Overall, the instance in Figure 1 admits a stable fractional matching with strictly greater welfare than any stable integral matching. □

Starting with the seminal work of Gale and Shapley [16], there is now an extensive literature on algorithms for computing stable solutions, including ones that optimize a variety of objectives pertaining to fairness and economic efficiency [17, 21, 25, 29, 36]. Most of these algorithms, however, are tailored to compute stable *integral* matchings. As Example 1 demonstrates, such algorithms could, in general, return highly suboptimal outcomes in our setting. Therefore, it becomes pertinent to understand the computational complexity of finding an "optimal" stable matching in the generalized model. Our work studies this question from the lens of the fundamental objective of social welfare, and asks the following natural question:

Can we efficiently compute an optimal or nearly optimal stable fractional matching?

1.1 Our results and roadmap

We formalize the above question by defining the optimization problem Optimal Stable Fractional Matching. To motivate this problem, we strengthen the observation in Example 1 to show that the social welfare gap between the best stable fractional and best stable integral matchings can be *arbitrarily* large. We show that the favorable welfare properties of stable fractional matchings come at the cost of limiting the algorithmic tools at our disposal. Specifically, we show that the set of stable fractional matchings can be non-convex, and that, in the worst case, stable fractional matchings can have a large support, thus prohibiting the use of support enumeration algorithms. Nevertheless, we present simple algorithms for Optimal Stable Fractional Matching with approximation ratio of $1 + \sigma_{\max}/\sigma_{\min}$, where $\sigma_{\max}$ and $\sigma_{\min}$ represent the maximum and minimum positive valuation in the input instance, respectively. For the variant Optimal ε-Stable Fractional Matching, where the stability constraints are relaxed by an ε factor, an embarrassingly simple algorithm computes $1/\varepsilon$-approximate solutions. We then proceed to our main result, which shows that these approximation guarantees are, somewhat surprisingly, almost the best achievable via polynomial-time algorithms (unless P = NP). To the best of our knowledge, these are the first computational complexity results for stable fractional matchings.

The rest of the paper is structured as follows. We present related work in the matching literature in Section 1.2. We continue in Section 2 with preliminary definitions and warm up with exponential-time algorithms that solve Optimal Stable Fractional Matching using linear programming. Section 3 presents the structural properties of (nearly)-optimal solutions of Optimal Stable Fractional Matching. Our algorithms are presented in Section 4, and our inapproximability results are presented in Section 5. All missing proofs can be found in the full version of the paper [7].

1.2 Further related work

The stable marriage problem has been extensively studied over the years [16, 18, 28, 34] with several interesting results in the original model of Gale and Shapley [16] as well as several of its variants [9, 20, 22, 24, 29]. The model of Gale and Shapley [16] admits efficient algorithms for finding stable matchings, but some of the above variants are computationally more challenging; we refer the interested reader to the survey of Iwama and Miyazaki [23] for a detailed exposition. Starting with the works of Vande Vate [37], Rothblum [35], and Roth et al. [32], there is by now a well-developed literature on the linear programming formulations of the stable matching problem [36, 38], and various combinatorial algorithms that optimize fairness-related objectives are also known [8, 15, 17, 21, 25]. However, the majority of this literature studies integral stable matchings under ordinal preferences.

The theoretical and practical importance of modeling agents' cardinal preferences in stable matching settings has been highlighted, among others, by Anshelevich et al. [4] and Pini et al. [30]. Anshelevich et al. [4] formulate the notions of exact/approximate stability in terms of cardinal preferences that are central to our work; see Definition 1. They study the "price of anarchy" for stable matchings (defined as the ratio of social welfare of the optimal matching and the worst stable integral matching) under various preference structures as well as its extensions to approximate stability. A similar approach has been adopted in related settings [2, 14]. Pini et al. [30] consider a notion of stability similar to [4], but focus on achieving economic efficiency (in particular, Pareto optimality and its variants) along with stability. They also study strategic aspects which are an exciting avenue for future research even in our model.

In terms of computing efficient (i.e., welfare-maximizing or cost-minimizing) stable matchings, Irving et al. [21], Manlove et al. [29], and Mai and Vazirani [26] provide efficient algorithms and/or inapproximability results, but crucially, these results apply to *integral* matchings only. To the best of our knowledge, the computational questions associated with computing stable *fractional* matchings have not been considered prior to our work.

Deligkas et al. [10] consider matchings computed by the natural greedy algorithm in edge-weighted graphs. Among other results, they show that the problem of computing the maximum-weight greedy matching in a bipartite graph is NP-hard. The greedy matchings in their model are stable integral matchings in ours. Still, as we discuss in the full version [7], this result implies the NP-hardness of Optimal Stable Fractional Matching. Our inapproximability results are much stronger though.

Our definition of fractional stability (in the presence of cardinal preferences) has appeared before in the economics literature [5, 11, 12, 19, 27]. However, these papers focus primarily on the relationship among various notions of stability and economic efficiency, and do not consider computational questions. We remark here that the term "fractional stability" has been overloaded in the related literature, as it has also been used to refer to fractional matchings that only have integral stable matchings in their support [36]; note that stability in the latter context can be defined purely in terms of the ordinal preferences. This notion of "ex-post" stability is fundamentally different from ours, as it is precisely the existence of unstable integral matchings in the support of the stable fractional matching (see Proposition 3) that allows for the large gain in welfare in Example 1 and Theorem 1. If one is interested in the ex-post notion of stability, then the computational problem clearly reduces to the case of integral matchings.

2 PRELIMINARIES

An instance of *Stable Matching problem with Cardinal preferences* (SMC) is given by the tuple $\langle M, W, U, V\rangle$, where $M := \{m_1, \ldots, m_n\}$ and $W := \{w_1, \ldots, w_n\}$ denote the set of n men and n women, respectively, and U and V are $n \times n$ matrices of non-negative rational numbers that specify the *valuations* of the agents. Specifically, $U(m, w)$ is the value derived by man m from his match with woman w, and $V(m, w)$ is the value derived by woman w from her match with man m. Many of our results will focus on two special classes of valuations, namely *binary* (where $U, V \in \{0, 1\}^{n\times n}$) and *ternary* valuations (where $U, V \in \{0, 1, \alpha\}^{n\times n}$ for some $\alpha > 1$).

We will often describe an SMC instance using its *graph representation*. An instance $\mathcal{I} = \langle M, W, U, V\rangle$ can be represented as a bipartite graph with vertex sets M and W, and an edge for every pair $(m, w) \in M \times W$ such that at least one of $U(m, w) > 0$ or $V(m, w) > 0$ holds. Each edge (m, w) in this graph has two valuations associated with it, namely $U(m, w)$ and $V(m, w)$.

A *fractional matching* $\mu : M \times W \to \mathbb{R}_{\geq 0}$ is an assignment of non-negative weights to all man-woman pairs such that $\sum_{w\in W} \mu(m, w) \leq 1$ for each $m \in M$ and $\sum_{m\in M} \mu(m, w) \leq 1$ for each $w \in W$. A fractional matching μ is said to be *complete* if $\sum_{w\in W} \mu(m, w) = 1$ for each man $m \in M$ and $\sum_{m\in M} \mu(m, w) = 1$ for each woman $w \in W$. An *integral* matching μ is a fractional matching with weights $\mu(m, w) \in \{0, 1\}$ for every pair (m, w). With slight abuse of notation, we sometimes view an integral matching μ as a set of pairs and write $(m, w) \in \mu$ instead of $\mu(m, w) = 1$.

It is well-known, and follows from the Birkhoff-von Neumann (BvN) theorem, that a fractional matching μ can be decomposed into a convex combination of $k = O(n^2)$ integral matchings $\mu^{(1)}, \mu^{(2)}, \ldots, \mu^{(k)}$ so that for every pair $(m, w) \in M \times W$, we have

$$\mu(m, w) = \textstyle\sum_{j=1}^{k} \lambda_j \cdot \mu^{(j)}(m, w),$$

where $\lambda_j > 0$ for all $j \in \{1, \ldots, k\}$ and $\sum_{j=1}^{k} \lambda_j = 1$. The set of integral matchings $\{\mu^{(1)}, \ldots, \mu^{(k)}\}$ is called the *support* of the fractional matching μ. Note that the support need not be unique.

We proceed with the formal definitions of *stability* and *approximate stability*, which in turn use the definition of the *utility* derived by agents under a fractional matching. In particular, the utility derived by the man m under μ is given by $u_m(\mu) := \sum_{w \in W} U(m,w)\mu(m,w)$, and the utility derived by the woman w is given by $v_w(\mu) := \sum_{m \in M} V(m,w)\mu(m,w)$.

Definition 1 (Stability). Given a fractional matching μ, a man-woman pair (m,w) is said to be a *blocking pair* if $u_m(\mu) < U(m,w)$ and $v_w(\mu) < V(m,w)$. A fractional matching μ is *stable* if there are no blocking pairs, i.e., for each $(m,w) \in M \times W$, either $u_m(\mu) \geq U(m,w)$ or $v_w(\mu) \geq V(m,w)$.

Definition 2 (ε-Stability). Given any $\varepsilon \in [0,1)$ and a fractional matching μ, a man-woman pair (m,w) is said to be ε*-blocking* if $u_m(\mu) < (1-\varepsilon)U(m,w)$ and $v_w(\mu) < (1-\varepsilon)V(m,w)$; otherwise, the pair is said to be ε*-stable*. A fractional matching μ is ε*-stable* if all pairs are ε-stable.

A stable fractional matching is also ε-stable for every $\varepsilon \geq 0$. The next statement follows from the seminal result of Gale and Shapley [16].

Proposition 1. *Given any* SMC *instance* $\mathcal{I}$*, a stable fractional matching* μ *for* $\mathcal{I}$ *always exists and can be computed in polynomial time.*

Proposition 1 was originally proven in [16] in the standard stable matching model with ordinal preferences and integral matchings. It is easy to see that given any SMC instance $\mathcal{I}$, if an integral matching μ is stable for an ordinal instance derived from $\mathcal{I}$ (where the ordinal preferences of each agent are consistent with its valuations, breaking ties arbitrarily), then it is also stable for the original instance $\mathcal{I}$.

Next, we define *social welfare*, which is a measure of the efficiency of a fractional matching.

Definition 3 (Social welfare). Given an SMC instance $\langle M, W, U, V\rangle$ and a fractional matching μ, the *social welfare* of μ is defined as

$$\mathcal{W}(\mu) := \sum_{m \in M} u_m(\mu) + \sum_{w \in W} v_w(\mu) = \sum_{m \in M} \sum_{w \in W} (U(m,w) + V(m,w))\mu(m,w).$$

An *optimal* matching is one with the highest social welfare among all fractional matchings. It follows from the BvN decomposition that there is always an integral optimal matching. Similarly, an *optimal stable* fractional matching (respectively, *optimal* ε*-stable* fractional matching) is one with the highest social welfare among all stable (respectively, all ε-stable) fractional matchings. We will use Optimal Stable Fractional Matching and Optimal ε-Stable Fractional Matching to refer to the corresponding optimization problems. For $\rho \in (0,1]$, the term ρ*-efficient* refers to a stable (respectively, ε-stable) fractional matching with welfare at least ρ times the welfare of the optimal stable (respectively, ε-stable) fractional matching. Thus, an optimal stable (or ε-stable) fractional matching is 1-efficient.

2.1 Computing optimal stable fractional matchings

We will now discuss two exponential-time algorithms for Optimal Stable Fractional Matching. The first algorithm uses the following mixed integer linear program (OPT-Stab):

$$\text{(OPT-Stab)} \quad \text{maximize} \quad \sum_{m \in M} u_m + \sum_{w \in W} v_w$$

$$\begin{aligned}
\text{subject to} \quad u_m &\geq U(m,w)y(m,w) && \forall\, m \in M, w \in W && (1)\\
v_w &\geq V(m,w)(1-y(m,w)) && \forall\, m \in M, w \in W && (2)\\
u_m &= \textstyle\sum_{w \in W} U(m,w)\mu(m,w) && \forall\, m \in M && (3)\\
v_w &= \textstyle\sum_{m \in M} V(m,w)\mu(m,w) && \forall\, w \in W && (4)\\
\textstyle\sum_{w \in W} \mu(m,w) &\leq 1 && \forall\, m \in M && (5)\\
\textstyle\sum_{m \in M} \mu(m,w) &\leq 1 && \forall\, w \in W && (6)\\
\mu(m,w) &\geq 0 && \forall\, m \in M, w \in W && (7)
\end{aligned}$$

$$y(m,w) \in \{0,1\} \qquad \forall\, m \in M, w \in W \tag{8}$$

The non-negative weights $\mu(m,w)$ of man-woman pairs as well as the utilities $u_m := u_m(\mu)$ and $v_w := v_w(\mu)$ of the agents (set in equalities (3) and (4)) are the fractional variables of (OPT-Stab). The binary variables $y(m,w)$ encode the stability requirements for pair (m,w) in constraints (1) and (2). Indeed, by setting $y(m,w)$ to 1 or 0, we can require either $u_m(\mu) \geq U(m,w)$ or $v_w(\mu) \geq V(m,w)$. Constraints (5) and (6) ensure feasibility. By enumerating over all possible combinations of values for the binary variables $y(m,w)$ for $(m,w) \in M \times W$, we get 2^{n^2} different linear programs, and, clearly, at least one of them must have the optimal stable fractional matching as its optimal solution.

Our second algorithm is slightly faster and solves at most $O(n^n)$ linear programs. It exploits the following linear program (OPT-Thresh), which is defined using non-negative constants θ_m for $m \in M$ and θ_w for $w \in W$, which we call *utility thresholds*.

(OPT-Thresh) maximize $\sum_{m\in M} u_m + \sum_{w\in W} v_w$

$$\begin{aligned}
\text{subject to}\quad & u_m \geq \theta_m && \forall\, m \in M && (9)\\
& v_w \geq \theta_w && \forall\, w \in W && (10)\\
& u_m = \textstyle\sum_{w\in W} U(m,w)\mu(m,w) && \forall\, m \in M && (11)\\
& v_w = \textstyle\sum_{m\in M} V(m,w)\mu(m,w) && \forall\, w \in W && (12)\\
& \textstyle\sum_{w\in W} \mu(m,w) \leq 1 && \forall\, m \in M && (13)\\
& \textstyle\sum_{m\in M} \mu(m,w) \leq 1 && \forall\, w \in W && (14)\\
& \mu(m,w) \geq 0 && \forall\, m \in M, w \in W && (15)
\end{aligned}$$

When all utility thresholds are set to zero, the solution of (OPT-Thresh) is an optimal (i.e., welfare-maximizing) fractional matching. Using (OPT-Thresh) to maximize social welfare under stability constraints is more challenging. We say that a set of utility thresholds is *stability-preserving* if for every pair of agents $m \in M$ and $w \in W$, either $\theta_m \geq U(m,w)$ or $\theta_w \geq V(m,w)$. Then, any fractional matching μ that is feasible for (OPT-Stab) is also feasible for (OPT-Thresh) for some stability-preserving set of utility thresholds. Conversely, any fractional matching μ that is feasible for (OPT-Thresh) with some set of stability-preserving utility thresholds is also feasible for (OPT-Stab). Therefore, in order to solve Optimal Stable Fractional Matching, it suffices to enumerate all $O(n^n)$ sets of utility thresholds with $\theta_m \in \{U(m,w) : w \in W\}$ for every man $m \in M$, compute utility thresholds $\theta_w \in \{V(m,w) : m \in M\}$ for all $w \in W$ so that the utility thresholds are stability-preserving, and solve (OPT-Thresh). Among these solutions, the fractional matching with highest social welfare will be the solution of Optimal Stable Fractional Matching.

3 STRUCTURAL PROPERTIES

In this section, we present several observations about the structure of optimal and nearly-optimal stable fractional matchings. We begin by considerably strengthening our observation in Example 1 regarding the welfare gap between stable fractional and stable integral matchings.

Theorem 1. *For every $\delta > 0$ and $\alpha \geq 2$, there exists an* SMC *instance with ternary valuations in $\{0,1,\alpha\}$ and an optimal stable fractional matching μ^* such that any stable integral matching μ^s satisfies $\mathcal{W}(\mu^s) \leq \left(\alpha - \frac{1}{2} - \delta\right)^{-1} \mathcal{W}(\mu^*)$.*

We emphasize that Theorem 1 is a *positive* result as it establishes that stable fractional matchings can have much higher welfare than their integral counterparts, and highlights the importance of Optimal Stable Fractional Matching. The proof of Theorem 1 appears in the full version [7].

Our next observation (Proposition 2) shows that the set of stable fractional matchings can be *non-convex* even for binary valuations. Interestingly, this does not prevent us from efficiently solving Optimal Stable Fractional Matching in this setting (see Theorem 4 in Section 4).

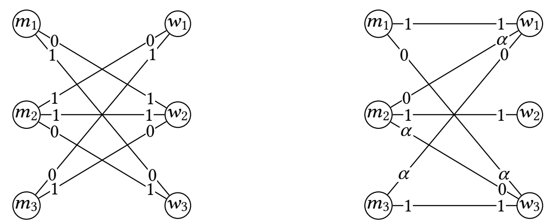

(a) **(b)**

Fig. 2. *The* SMC *instances used in the proofs of Propositions 2 and 3.*

Proposition 2. *There exists an* SMC *instance with binary valuations for which the set of stable fractional matchings is non-convex.*

Proof. Consider the instance $\mathcal{I} = \langle M, W, U, V \rangle$ with three men m_1, m_2, m_3 and three women w_1, w_2, w_3, whose graph representation and agent valuations are shown in Figure 2a. Consider the integral matchings $\mu^{(1)} := \{(m_1, w_3), (m_2, w_1), (m_3, w_2)\}$ and $\mu^{(2)} := \{(m_1, w_2), (m_2, w_3), (m_3, w_1)\}$. It is easy to verify that both $\mu^{(1)}$ and $\mu^{(2)}$ are stable for $\mathcal{I}$. However, the fractional matching $\mu := 0.5\mu^{(1)} + 0.5\mu^{(2)}$ is not stable since (m_2, w_2) is a blocking pair; indeed, $0.5 = u_{m_2}(\mu) < U(m_2, w_2) = 1$ and $0.5 = v_{w_2}(\mu) < V(m_2, w_2) = 1$. □

The structure of stable fractional matchings becomes much more interesting (and, as we will see in Section 5, also computationally unwieldy) when we move to *ternary* valuations. It turns out that the support of a stable fractional matching can comprise entirely of *unstable* integral matchings (Proposition 3), and its size can grow *linearly* with the input (Theorem 3). These observations pose major limitations on the set of algorithmic tools at our disposal.

Proposition 3. *There exists an* SMC *instance with ternary valuations and a stable fractional matching μ such that every integral matching in any support of μ is unstable.*

Proof. Consider the SMC instance $\mathcal{I} = \langle M, W, U, V \rangle$ with three men and three women shown in Figure 2b. The parameter $\alpha \geq 3$ is a constant. There are six different perfect integral matchings:

- Matching $\mu^{(1)}$, which consists of pairs (m_1, w_1), (m_2, w_2), and (m_3, w_3) and has a social welfare of 6. It is easy to verify that this is the unique stable integral matching. Also, any subset of $\mu^{(1)}$ is not stable as the pair that is missing from $\mu^{(1)}$ will be blocking.
- Matching $\mu^{(2)}$, which consists of pairs (m_1, w_2), (m_2, w_3), and (m_3, w_1) and has a social welfare of 2α. The matching is not stable since the pair (m_1, w_1) is blocking.
- Matching $\mu^{(3)}$, which consists of pairs (m_1, w_3), (m_2, w_1), and (m_3, w_2) and has a social welfare of 2α. It is not stable since (m_2, w_2) is blocking.
- Matching $\mu^{(4)}$, which consists of pairs (m_1, w_1), (m_2, w_3), and (m_3, w_2) and has a social welfare of $\alpha + 2$. It is not stable since (m_3, w_3) is blocking.
- Matching $\mu^{(5)}$, which consists of pairs (m_1, w_3), (m_2, w_2), and (m_3, w_1) and has a social welfare of $2\alpha + 2$. It is not stable since (m_1, w_1) is blocking.
- Matching $\mu^{(6)}$, which consists of pairs (m_1, w_2), (m_2, w_1), and (m_3, w_3) and has a social welfare of $\alpha + 2$. It is not stable since the pair (m_2, w_2) is blocking.

Consider the matching $\mu := \frac{1}{\alpha(\alpha-1)} \cdot \mu^{(2)} + \frac{1}{\alpha} \cdot \mu^{(3)} + \frac{\alpha-2}{\alpha-1} \cdot \mu^{(5)}$. It is easy to verify that μ is stable. Indeed, the utilities of the agents under μ are given by $u_{m_1}(\mu) = 0$, $u_{m_2}(\mu) = 1$, $u_{m_3}(\mu) = \alpha - 1$, $v_{w_1}(\mu) = 1$, $v_{w_2}(\mu) = \frac{\alpha-2}{\alpha-1}$ and $v_{w_3}(\mu) = \alpha - \frac{1}{\alpha-1}$. Notice that only the pairs (m_1, w_1), (m_2, w_2), and (m_3, w_3) need to be checked for stability, since any other pair has at least one agent with a valuation of zero (and, hence, the stability constraint for those pairs is trivially satisfied). For each of the pairs (m_1, w_1), (m_2, w_2), and (m_3, w_3), there is some member of the pair that has a utility of at least 1 under μ, which meets the requisite stability threshold, implying that μ is stable. Finally, notice that $\mu(m_1, w_1) = 0$, which means that the unique stable integral matching $\mu^{(1)}$ cannot occur in a support of μ. □

We remark that with some extra work, one can show that the matching μ in the proof of Proposition 3 is the unique optimal stable fractional matching.

As mentioned previously in Section 2, a (stable) fractional matching is the convex combination of at most n^2 integral ones. Theorem 2 provides a stronger bound on the support size of an *optimal* stable fractional matching.

Theorem 2. *Given any* SMC *instance $\mathcal{I}$, there exists an optimal stable fractional matching for $\mathcal{I}$ with at most $4n$ integral matchings in its support.*

Proof. Let μ^* be an optimal stable fractional matching for $\mathcal{I}$. Recall from Section 2.1 that μ^* solves the program (OPT-Thresh) for some set of stability-preserving utility thresholds. Observe that (OPT-Thresh) has n^2 free variables (we ignore here the $2n$ variables u_m for $m \in M$ and v_w for $w \in W$, which depend on the remaining ones according to constraints (11) and (12)). Without loss of generality, μ^* is an optimal *extreme point* solution of (OPT-Thresh). That is, when (OPT-Thresh) is instantiated for μ^*, n^2 linearly independent inequality constraints become tight. Among them, at most $4n$ can correspond to the sets of constraints (9), (10), (13), and (14). The remaining ones must correspond to the set of constraints (15), implying that at least $n^2 - 4n$ free variables will be equal to zero. Thus, μ^* can assign positive weights to at most $4n$ man-woman pairs and, consequently, can have at most $4n$ integral matchings in its support. □

Next we show that the bound in Theorem 2 is tight up to a constant factor.

Theorem 3. *For every $\rho \in (0, 1]$, there exists a family of* SMC *instances with ternary valuations for which any support of a ρ-efficient stable fractional matching consists of $\Omega(\rho n)$ integral matchings.*

Proof. Consider a family of SMC instances $\mathcal{I}_n = \langle M, W, U, V\rangle$ with $M = \{m_1, \ldots, m_n\}$ and $W = \{w_1, \ldots, w_n\}$, where n is odd. Let α be such that $\alpha > \max\left\{n + 2, \frac{2n}{\rho(n-1)}\right\}$. The (ternary) valuations of the agents are defined as follows: For each $i \in \{1, 2, \ldots, n\}$, $U(m_i, w_i) = V(m_i, w_i) = 1$. For each $i \in \{1, 2, \ldots, \frac{n-1}{2}\}$, $U(m_{2i}, w_1) = U(m_{2i+1}, w_{2i}) = V(m_{2i}, w_{2i+1}) = \alpha$ and $V(m_{2i}, w_1) = V(m_{2i+1}, w_{2i}) = U(m_{2i}, w_{2i+1}) = 0$. Finally, $U(m_n, w_1) = 0$ and $V(m_n, w_1) = \alpha$. For all remaining pairs $(m, w) \in M \times W$, $U(m, w) = V(m, w) = 0$. Figure 3a illustrates the SMC instance $\mathcal{I}_5$.

Define $\mu^{\text{opt}} := \{(m_1, w_1)\} \cup \left\{(m_{2i}, w_{2i+1}), (m_{2i+1}, w_{2i}) : i \in \{1, 2, \ldots, \frac{n-1}{2}\}\right\}$ (see Figure 3b). We also define a number of other integral matchings obtained by modifying μ^{opt}, as follows: For $i \in \{1, 2, \ldots, \frac{n-1}{2}\}$, the matching $\mu^{(i)}$ (see Figures 3c and 3d) is the integral matching which is obtained from μ^{opt} by replacing $\{(m_1, w_1), (m_{2i}, w_{2i+1})\}$ with $\{(m_1, w_{2i+1}), (m_{2i}, w_1)\}$, i.e.,

$$\mu^{(i)} := (m_1, w_{2i+1}) \cup (m_{2i}, w_1) \cup (m_{2i+1}, w_{2i}) \cup \{(m_{2\ell}, w_{2\ell+1}) \cup (m_{2\ell+1}, w_{2\ell})\}_{\ell \in \{1,2,\ldots,\frac{n-1}{2}\} \setminus \{i\}}\,.$$

Also, the matching $\mu^{(\frac{n+1}{2})}$ (see Figure 3e) is the integral matching obtained from μ^{opt} by replacing $\{(m_1, w_1), (m_n, w_{n-1})\}$ with $\{(m_1, w_{n-1}), (m_n, w_1)\}$, i.e.,

$$\mu^{(\frac{n+1}{2})} := (m_1, w_{n-1}) \cup (m_{n-1}, w_n) \cup (m_n, w_1) \cup \{(m_{2\ell}, w_{2\ell+1}) \cup (m_{2\ell+1}, w_{2\ell})\}_{\ell \in \{1,2,\ldots,\frac{n-3}{2}\}}\,.$$

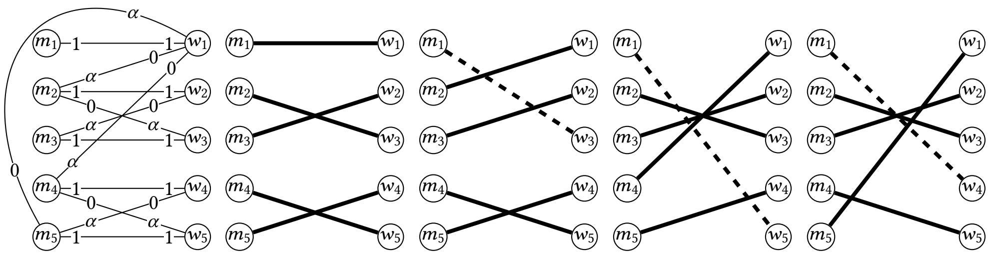

(a) **(b)** **(c)** **(d)** **(e)**

Fig. 3. *Subfigure (a) illustrates the graph representation of the* SMC *instance $\mathcal{I}_n$ described in the proof of Theorem 3 for $n = 5$. Subfigure (b) shows the matching μ^{opt}. Subfigures (c), (d), and (e) show the matchings $\mu^{(1)}$, $\mu^{(2)}$, and $\mu^{(3)}$, respectively. Dashed lines indicate zero-valuation pairs that do not appear in the graph representation.*

Now, consider the fractional matching $\mu := \sum_{i=1}^{\frac{n+1}{2}} \frac{1}{\alpha}\mu^{(i)} + \left(1 - \frac{n+1}{2\alpha}\right)\mu^{\text{opt}}$. Since $\alpha > n+2 > \frac{n+1}{2}$, μ is well-defined and has the matchings μ^{opt} and $\mu^{(i)}$ for $i \in \left\{1, \ldots, \frac{n+1}{2}\right\}$ in its support. Notice that $\mathcal{W}(\mu^{\text{opt}}) = (n-1)\alpha + 2$ and $\mathcal{W}(\mu^{(i)}) = (n-1)\alpha$ for all $i \in \{1, 2, \ldots, \frac{n+1}{2}\}$. Thus, $\mathcal{W}(\mu) > (n-1)\alpha$.

It can be verified that μ is stable. Indeed, we only need to check the blocking condition for the pairs (m_i, w_i) with $i \in \{1, 2, \ldots, n\}$. We have that $v_{w_1}(\mu) \geq 1$ (since $\mu^{(\frac{n+1}{2})}$ has weight $\frac{1}{\alpha}$ in μ and $V(m_n, w_1) = \alpha$), $u_{m_{2i}}(\mu) \geq 1$ for each $i \in \{1, 2, \ldots, \frac{n-1}{2}\}$ (since $\mu^{(i)}$ has weight $\frac{1}{\alpha}$ in μ and $U(m_{2i}, w_1) = \alpha$), and $v_{w_{2i+1}}(\mu) \geq 1$ for each $i \in \{1, 2, \ldots, \frac{n-1}{2}\}$ (since $\mu^{(i)}$ has weight $\frac{1}{\alpha}$ in μ and $V(m_{2i}, w_{2i+1}) = \alpha$). The welfare of the optimal stable fractional matching must therefore be at least $\mathcal{W}(\mu)$, and thus strictly greater than $(n-1)\alpha$.

We now claim that any ρ-efficient stable fractional matching μ' satisfies $\mu'(m_n, w_1) > 0$. Indeed, assuming otherwise that $\mu'(m_n, w_1) = 0$, the only pair that can give positive utility to man m_1 and woman w_1 is (m_1, w_1). Hence, we must also have $\mu'(m_1, w_1) = 1$, and, as a result, $\mu'(m_{2i}, w_1) = 0$ for $i \in \left\{1, 2, \ldots, \frac{n-1}{2}\right\}$. Then, the only pair that can give positive utility to man m_{2i} and woman w_{2i} is (m_{2i}, w_{2i}), and hence, it must also be that $\mu'(m_{2i}, w_{2i}) = 1$. Consequently, the only pair that can give positive utility to man m_{2i+1} and woman w_{2i+1} is (m_{2i+1}, w_{2i+1}) and, hence, we must have $\mu'(m_{2i+1}, w_{2i+1}) = 1$. The welfare of matching μ' would then be $2n$, which is less than $\rho(n-1)\alpha$ by the assumed bound on α. In other words, the welfare of μ' would be less than ρ times the welfare of the stable fractional matching μ, contradicting the assumption that μ' is ρ-efficient.

The final step in the proof involves showing that for any stable fractional matching μ' with support of size at most $\frac{n-1}{2}\rho$, we must have $\mathcal{W}(\mu') < \rho(n-1)\alpha$; the desired bound on the support size would then follow from the contrapositive. Let $T := \left\{i \in \left\{1, 2, \ldots \frac{n-1}{2}\right\} : \mu'(m_{2i}, w_1) > 0\right\}$, and $\overline{T} := \left\{1, 2, \ldots \frac{n-1}{2}\right\} \setminus T$. Since μ' has support of size at most $\frac{n-1}{2}\rho$ and $\mu'(m_n, w_1) > 0$, it holds that $|T| \leq \frac{n-1}{2}\rho - 1$.

For every $i \in T$, the agents m_{2i}, w_{2i}, m_{2i+1}, and w_{2i+1} can together contribute at most 2α to the welfare. On the other hand, when $i \in \overline{T}$, we have $\mu'(m_{2i}, w_1) = 0$, and the only pair that can give positive utility to man m_{2i} and woman w_{2i} is (m_{2i}, w_{2i}). Therefore, we must have that $\mu'(m_{2i}, w_{2i}) = 1$. Consequently, the only pair that can give positive utility to man m_{2i+1} and woman w_{2i+1} is (m_{2i+1}, w_{2i+1}), and it follows that $\mu'(m_{2i+1}, w_{2i+1}) = 1$. Therefore, when $i \in \overline{T}$, the agents m_{2i}, w_{2i}, m_{2i+1}, and w_{2i+1} can together contribute at most 4 to the welfare. Taking the possible contribution of pair (m_1, w_1) into account, we have that

$$\mathcal{W}(\mu') \leq 2 + 2\alpha|T| + 4|\overline{T}| = 2n + (2\alpha - 4)|T| \leq 2n + \rho(n-1)\alpha - 2\alpha - 2(n-1)\rho + 4 < \rho(n-1)\alpha.$$

The equality follows from the definition of $\overline{T}$, the second inequality follows from the bound on $|T|$ above, and the third one from the definition of α. This completes the proof of Theorem 3. □

Theorem 3 has an interesting algorithmic implication. As the support size can be large in any good approximation of the optimal stable fractional matching, it implies that *support enumeration* strategies—which have been proved useful in other contexts; see [6] and the references therein—will be ineffective in computing (even approximate) solutions of Optimal Stable Fractional Matching. A similar implication can be shown for optimal ε-stable fractional matchings. In contrast, as we will show in Section 4, ε-stable fractional matchings of *small* support can be easily computed, and provide nearly the best approximation ratios achievable by efficient algorithms (under standard complexity-theoretic assumptions).

4 ALGORITHMIC RESULTS

We begin the discussion of our algorithmic results with binary valuations. In this setting, Optimal Stable Fractional Matching reduces to computing a maximum weight matching on a specific weighted graph associated with the given instance.

Theorem 4. *Given an* SMC *instance* $\mathcal{I} = \langle M, W, U, V \rangle$ *with binary valuations, an optimal stable fractional matching for* $\mathcal{I}$ *can be computed in polynomial time.*

Next, we consider general valuations and show how to exploit stable integral matchings to get an approximate solution for Optimal Stable Fractional Matching. Let $\sigma_{\max}$ and $\sigma_{\min}$ denote the largest and the smallest non-zero valuation among all agents in $\mathcal{I}$, respectively. We call a man-woman pair (m, w) *light* if either $U(m, w) = 0$ or $V(m, w) = 0$, and *heavy* otherwise. Given an SMC instance $\mathcal{I}$ as input, our algorithm computes a stable integral matching for $\mathcal{I}$, say μ, in two steps: First, it computes a stable integral matching μ_1 using only the heavy pairs (and taking into account the stability constraints in heavy pairs only). Then, it *completes* the solution with a matching μ_2 of maximum welfare using the light pairs subject to feasibility constraints, i.e., using light pairs that do not share any agents with the pairs in μ_1. The light pairs impose no additional constraints on stability, so the resulting matching is stable.

We will show that μ has approximation ratio $1 + \sigma_{\max}/\sigma_{\min}$. Let μ^{opt} be an optimal matching for $\mathcal{I}$. Also, let μ_1^{opt} be the set of pairs of μ^{opt} that share an agent with some pair of μ_1, i.e., $\mu_1^{\text{opt}} := \{(m, w) \in \mu^{\text{opt}} : \text{at least one of } m \text{ or } w \text{ is matched under } \mu_1\}$. By definition of μ_2, we have $\mathcal{W}(\mu_2) \geq \mathcal{W}(\mu^{\text{opt}} \setminus \mu_1^{\text{opt}})$. To complete the proof, we will need the following lemma.

Lemma 1. $\mathcal{W}(\mu_1 \setminus \mu_1^{opt}) \geq (1 + \sigma_{\max}/\sigma_{\min})^{-1} \mathcal{W}(\mu_1^{opt} \setminus \mu_1)$.

Proof. Our proof constructs a mapping in which every pair $(m, w) \in \mu_1^{\text{opt}} \setminus \mu_1$ is mapped to one of its agents, whom we will call the *witness* of the pair. The mapping is such that the utility of the witness in the matching $\mu_1 \setminus \mu_1^{\text{opt}}$ is at least $(1 + \sigma_{\max}/\sigma_{\min})^{-1} (U(m, w) + V(m, w))$. Note that once we establish the said mapping, the proof will follow, since each agent can be the witness of at most one pair of $\mu_1^{\text{opt}} \setminus \mu_1$ and $\mathcal{W}(\mu_1 \setminus \mu_1^{\text{opt}})$ is at least the total utility of the witnesses in $\mu_1 \setminus \mu_1^{\text{opt}}$.

Consider a light pair $(m, w) \in \mu_1^{\text{opt}} \setminus \mu_1$. The witness is an agent (m or w) who also belongs to a pair of $\mu_1 \setminus \mu_1^{\text{opt}}$; such an agent certainly exists by the definition of μ_1^{opt}. Since all pairs of $\mu_1 \setminus \mu_1^{\text{opt}}$ are heavy, the utility of the witness of (m, w) in $\mu_1 \setminus \mu_1^{\text{opt}}$ is at least $\sigma_{\min} = \frac{\sigma_{\min}}{\sigma_{\max}}(0 + \sigma_{\max}) \geq (1 + \sigma_{\max}/\sigma_{\min})^{-1} (U(m, w) + V(m, w))$, since (m, w) is light.

Now consider a heavy pair $(m, w) \in \mu_1^{\text{opt}} \backslash \mu_1$. If μ_1 contains a pair (m, w') with $U(m, w') \geq U(m, w)$, select agent m to be the witness, otherwise select agent w. Note that in the latter case, stability of μ_1 implies the existence of $(m', w) \in \mu_1$ such that $V(m', w) \geq V(m, w)$. Hence, the utility of the witness of (m, w) in $\mu_1 \backslash \mu_1^{\text{opt}}$ is at least $\min\{U(m, w), V(m, w)\} \geq (1 + \sigma_{\max}/\sigma_{\min})^{-1} (U(m, w) + V(m, w))$. □

Now, Lemma 1 gives the desired approximation ratio, as follows:

$$\begin{aligned} \mathcal{W}(\mu) &= \mathcal{W}(\mu_1) + \mathcal{W}(\mu_2) = \mathcal{W}(\mu_1 \setminus \mu_1^{\text{opt}}) + \mathcal{W}(\mu_1 \cap \mu_1^{\text{opt}}) + \mathcal{W}(\mu_2) \\ &\geq (1 + \sigma_{\max}/\sigma_{\min})^{-1}\, \mathcal{W}(\mu_1^{\text{opt}} \setminus \mu_1) + \mathcal{W}(\mu_1^{\text{opt}} \cap \mu_1) + \mathcal{W}(\mu^{\text{opt}} \setminus \mu_1^{\text{opt}}) \\ &\geq (1 + \sigma_{\max}/\sigma_{\min})^{-1}\, \mathcal{W}(\mu^{\text{opt}}). \end{aligned}$$

For ternary valuations in $\{0, 1, \alpha\}$, the above algorithm gives a $(1 + \alpha)$-approximation. An improved approximation for ternary valuations can be achieved using the following modification: When computing the stable integral matching, resolve ties in favour of the pairs (m, w) with the highest $U(m, w) + V(m, w)$. The next lemma establishes an improved approximation ratio of $\max\{2, \alpha\}$.

Lemma 2. *The modified algorithm for* SMC *instances with ternary valuations in* $\{0, 1, \alpha\}$ *satisfies* $\mathcal{W}(\mu_1 \setminus \mu_1^{opt}) \geq \min\{\frac{1}{2}, \frac{1}{\alpha}\}\mathcal{W}(\mu_1^{opt} \setminus \mu_1)$.

Proof. For a pair (m, w) of matching $\mu_1^{\text{opt}} \setminus \mu_1$, we use the term *neighborhood* to refer to the pairs of $\mu_1 \setminus \mu_1^{\text{opt}}$ that use agent m or w. We will show that the total utility from pairs in the neighborhood of (m, w) is at least $\min\{1, 2/\alpha\} (U(m, w) + V(m, w))$. Since each pair of $\mu_1 \setminus \mu_1^{\text{opt}}$ can be in the neighborhood of at most two pairs of $\mu_1^{\text{opt}} \setminus \mu_1$, this will give us the desired inequality.

Indeed, by the particular way we resolve ties in the ordinal preferences before computing the matching μ_1, a heavy pair (m, w) in $\mu_1^{\text{opt}} \setminus \mu_1$ must have a pair of utility at least $U(m, w) + V(m, w)$ in its neighborhood. A light pair (m, w) has $U(m, w) + V(m, w) \leq \alpha$ and certainly has a heavy pair of utility at least 2 in its neighborhood. □

The above discussion is summarized in the following statement.

Theorem 5. *There is a polynomial-time algorithm which, given an* SMC *instance* $\mathcal{I}$ *with an optimal matching* μ^{opt}*, computes a stable integral matching* μ *with* $\mathcal{W}(\mu) \geq \min\{\frac{1}{2}, \frac{1}{\alpha}\}\mathcal{W}(\mu^{opt})$ *if* $\mathcal{I}$ *has ternary valuations in* $\{0, 1, \alpha\}$*, and* $\mathcal{W}(\mu) \geq (1 + \sigma_{\max}/\sigma_{\min})^{-1}\, \mathcal{W}(\mu^{opt})$ *in general, where* $\sigma_{\max}$ *and* $\sigma_{\min}$ *denote the highest and the lowest non-zero valuation in* $\mathcal{I}$*, respectively.*

We conclude this section by considering approximate stability. For general valuations, we present a polynomial-time $1/\varepsilon$-approximation algorithm for Optimal ε-Stable Fractional Matching, which constructs an ε-stable fractional matching with a small support by combining an optimal matching with a stable integral matching.

Theorem 6. *There is a polynomial-time algorithm that given any* SMC *instance* $\mathcal{I} = \langle M, W, U, V\rangle$ *and any rational* $\varepsilon \in [0, 1]$*, computes a fractional matching* μ *that is* ε*-stable for* $\mathcal{I}$ *such that* $\mathcal{W}(\mu) \geq \varepsilon\mathcal{W}(\mu^{opt})$*, where* μ^{opt} *is an optimal matching for* $\mathcal{I}$*.*

Proof. Let μ^{s} be any stable integral matching and μ^{opt} be an optimal matching for $\mathcal{I}$. Note that both μ^{s} and μ^{opt} can be computed in polynomial time. We will show that $\mu := (1 - \varepsilon)\mu^{\text{s}} + \varepsilon\mu^{\text{opt}}$ satisfies the desired properties. Indeed, $\mathcal{W}(\mu) = (1-\varepsilon)\mathcal{W}(\mu^{\text{s}}) + \varepsilon\mathcal{W}(\mu^{\text{opt}}) \geq \varepsilon\mathcal{W}(\mu^{\text{opt}})$. Furthermore, since μ^{s} is stable, we have that for any man-woman pair $(m, w) \in M \times W$, either $u_m(\mu^{\text{s}}) \geq U(m, w)$ or $v_w(\mu^{\text{s}}) \geq V(m, w)$. The former condition implies that $u_m(\mu) \geq (1 - \varepsilon)u_m(\mu^{\text{s}}) \geq (1 - \varepsilon)U(m, w)$, while the latter condition gives $v_w(\mu) \geq (1 - \varepsilon)V(m, w)$. Either way, the pair (m, w) is ε-stable. □

5 HARDNESS OF APPROXIMATION

In this section, we present our inapproximability statements, which are by far the technically most involved results in the paper. We present polynomial-time reductions which, given a 3SAT formula ϕ of a particular structure, construct SMC instances that simulate the evaluation of ϕ for every variable assignment. The constructed SMC instances consist of several gadgets including an *accumulator*. The simulation of the evaluation of ϕ by the SMC instance is such that:

(a) when ϕ has a satisfying assignment, there is a stable (or ε-stable) fractional matching where the contribution of the agents in the accumulator gadget to the welfare can be large and dominates the contribution from the remaining SMC instance and
(b) when ϕ is not satisfiable, the contribution of the accumulator and, subsequently, the total welfare of any stable (or ε-stable) fractional matching is very small.

Hence, distinguishing between SMC instances with stable (or ε-stable) fractional matchings of very high and very low welfare would allow us to decide 3SAT. We have two inapproximability statements: Theorem 7 for Optimal Stable Fractional Matching and Theorem 8 for Optimal ε-Stable Fractional Matching.

Theorem 7. *For every constant $\delta > 0$, it is NP-hard to approximate* Optimal Stable Fractional Matching *for* SMC *instances with ternary valuations in* $\{0, 1, \alpha\}$ *to within a factor of (i)* $\alpha - 1/2 - \delta$ *if* $\alpha = O(n)$*, and (ii)* $\Omega(n^{1-\delta})$ *otherwise.*

Theorem 8. *For any constants* $\varepsilon \in (0, 0.03]$ *and* $\delta > 0$*, it is NP-hard to approximate* Optimal ε-Stable Fractional Matching *to within a factor of* $1/\varepsilon - \delta$.

We will prove Theorem 7 here; the proof of Theorem 8, which uses similar gadgets but is slightly more involved, appears in the full version [7]. Since the proof is long, we have divided it into three parts: the description of the reduction (Section 5.1), technical claims with gadget properties (Section 5.2), and the proof of the inapproximability result (Section 5.3).

5.1 The reduction

In particular, we present a polynomial-time reduction from 2P2N-3SAT, the special case of 3SAT consisting of 3-CNF clauses in which every variable appears four times: twice as a positive literal and twice as a negative one. 2P2N-3SAT is known to be NP-hard [39]. Our reduction takes as input an instance of 2P2N-3SAT consisting of N (boolean) variables $x_1, x_2, ..., x_N$, and a 3-CNF formula ϕ with $L = 4N/3$ clauses $c_1, c_2, ..., c_L$. Without loss of generality, we assume that each clause in ϕ consists of distinct literals.

Given the instance of 2P2N-3SAT, our reduction generates an instance $\mathcal{I} = \langle M, W, U, V \rangle$ of Optimal Stable Fractional Matching. As usual, we denote by n the number of men (or women) in $\mathcal{I}$. We will use a positive integer parameter k which will determine the size of n; in particular, $n = O(N+k)$. We define $\mathcal{I}$ by referring to its graph representation, which consists of *variable gadgets*, *clause gadgets*, *variable-clause connectors*, an *accumulator*, and *clause-accumulator connectors*. For each gadget, we classify the edges (i.e., man-woman pairs and their valuations) into the following three types:

- *man-heavy* edges (m, w) with $U(m, w) = \alpha$ and $V(m, w) = 0$,
- *woman-heavy* edges (m, w) with $U(m, w) = 0$ and $V(m, w) = \alpha$, and
- *balanced* edges (m, w) with $U(m, w) = V(m, w) = 1$.

Recall that any pair (m, w) that does not appear as an edge in the graph representation has $U(m, w) = V(m, w) = 0$.

The instance $\mathcal{I}$ has a variable gadget for every variable x, which consists of five men m_1^x, m_2^x, e_1^x, e_2^x, e_3^x, four women $w_1^{\overline{x}}$, $w_2^{\overline{x}}$, f_1^x, f_2^x and the ten balanced edges (e_1^x, f_1^x), (m_1^x, f_1^x), $(m_1^x, w_1^{\overline{x}})$, $(e_3^x, w_1^{\overline{x}})$, $(e_3^x, w_2^{\overline{x}})$, $(m_2^x, w_2^{\overline{x}})$, (m_2^x, f_2^x), (e_2^x, f_2^x), $(m_1^x, w_2^{\overline{x}})$, and $(m_2^x, w_1^{\overline{x}})$, as shown in Figure 4a.

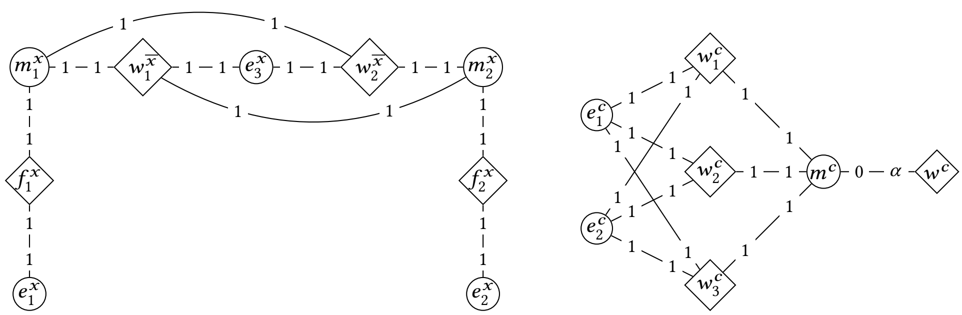

(a) **(b)**

Fig. 4. *(a) The variable gadget corresponding to the variable x. (b) The clause gadget corresponding to the clause c and its CA-connector (m^c, w^c). As a convention, we use circles to represent men and diamonds to represent women.*

For every clause c, instance $\mathcal{I}$ has a clause gadget with three men m^c, e_1^c, e_2^c, three women w_1^c, w_2^c, w_3^c, and the nine balanced edges between them, as shown in Figure 4b.

For every appearance of a literal in a clause, there is a variable-clause connector (or *VC-connector*). VC-connectors have different structure depending (1) on whether they correspond to positive or negative literals, and (2) on the value of α. In each case, we identify one edge of the VC-connector as the *input*, and either one or two edges as the *output*.

Specifically, for every positive literal x whose i-th appearance ($i \in \{1, 2\}$) is as the j-th literal ($j \in \{1, 2, 3\}$) of clause c, $\mathcal{I}$ has a VC-connector defined as follows:

- When $\alpha \geq 2$, the VC-connector consists of a single woman-heavy edge between m_i^x (from the variable gadget corresponding to variable x) and w_j^c (from the clause gadget corresponding to clause c), as shown in Figure 5a. This edge is simultaneously the input and the output edge of the VC-connector.
- When $\alpha \in (3/2, 2)$, the VC-connector consists of woman $w^{x,c}$, man $m^{x,c}$, the woman-heavy edges $(m_i^x, w^{x,c})$ and $(m^{x,c}, w_j^c)$, and the balanced edge $(m^{x,c}, w^{x,c})$, as shown in Figure 5b. Here, $(m_i^x, w^{x,c})$ is the input and $(m^{x,c}, w_j^c)$ is the output edge.

For every negative literal $\overline{x}$ whose i-th appearance ($i \in \{1, 2\}$) is as the j-th literal ($j \in \{1, 2, 3\}$) of clause c, $\mathcal{I}$ has a VC-connector defined as follows:

- When $\alpha \geq 2$, the VC-connector consists of man $m^{\overline{x},c}$, woman $w^{\overline{x},c}$, the man-heavy edge $(m^{\overline{x},c}, w_i^{\overline{x}})$, the balanced edge $(m^{\overline{x},c}, w^{\overline{x},c})$, and the woman-heavy edge $(m^{\overline{x},c}, w_j^c)$, as shown in Figure 5c. Here, $(m^{\overline{x},c}, w_i^{\overline{x}})$ is the input and $(m^{\overline{x},c}, w_j^c)$ is the output edge.
- When $\alpha \in (3/2, 2)$, the VC-connector consists of three men $m_1^{\overline{x},c}$, $m_2^{\overline{x},c}$, $m_3^{\overline{x},c}$, three women $w_1^{\overline{x},c}$, $w_2^{\overline{x},c}$, $w_3^{\overline{x},c}$, the man-heavy edges $(m_1^{\overline{x},c}, w_i^{\overline{x}})$, $(m_3^{\overline{x},c}, w_1^{\overline{x},c})$, $(m_3^{\overline{x},c}, w_2^{\overline{x},c})$, the woman-heavy edges $(m_1^{\overline{x},c}, w_2^{\overline{x},c})$, $(m_2^{\overline{x},c}, w_j^c)$, $(m_3^{\overline{x},c}, w_j^c)$, and the balanced edges $(m_1^{\overline{x},c}, w_1^{\overline{x},c})$, $(m_2^{\overline{x},c}, w_2^{\overline{x},c})$, $(m_3^{\overline{x},c}, w_3^{\overline{x},c})$, as shown in Figure 5d. In this case, the VC-connector has one input edge $(m_1^{\overline{x},c}, w_i^{\overline{x}})$ and two output edges $(m_2^{\overline{x},c}, w_j^c)$ and $(m_3^{\overline{x},c}, w_j^c)$.

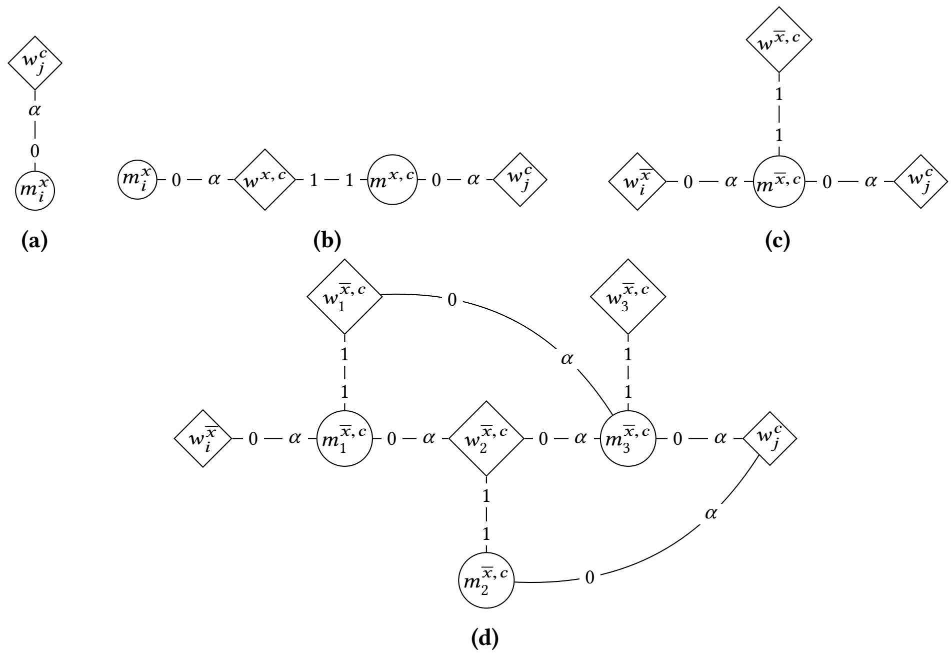

Fig. 5. *VC-connectors corresponding to clause c and positive literal x for (a) $\alpha \geq 2$ and (b) $\alpha \in (3/2, 2)$, and to clause c and negative literal $\overline{x}$ for (c) $\alpha \geq 2$ and (d) $\alpha \in (3/2, 2)$.*

The accumulator (Figure 6) of instance $\mathcal{I}$ has different structure depending on the value of α. Its size depends on the positive integer parameter k.

- When $\alpha \geq 2$ (see Figure 6a), the accumulator has man m_i and woman w_i for all $i \in \{1, \ldots, k\}$, men e_i^1 and e_i^2 and woman f_i^1 for all $i \in \{1, \ldots, k-1\}$, man e_i^3 and women f_i^2 and f_i^3 for all $i \in \{2, \ldots, k\}$, and woman w^c for every clause c of ϕ. In addition, there are man-heavy edges (m_i, w_{i-1}) and (e_i^3, f_i^2) for all $i \in \{2, \ldots, k\}$ and (e_i^2, w_i) for all $i \in \{1, \ldots, k-1\}$, the balanced edges (m_1, w^c) for every clause c, which we call *tine* edges, (e_i^1, w_i) for all $i \in \{1, \ldots, k-1\}$ and (m_i, f_i^2) for all $i \in \{2, \ldots, k\}$, and the woman-heavy edges (m_i, w_i) for all $i \in \{1, \ldots, k\}$, (e_i^1, f_i^1) for all $i \in \{1, \ldots, k-1\}$, and (m_i, f_i^3) for all $i \in \{2, \ldots, k\}$.
- When $\alpha \in (3/2, 2)$ (see Figure 6b), the accumulator has man m_i, woman w_i for $i \in \{1, \ldots, k\}$, man e_i^1 and woman f_i^1 for $i \in \{1, \ldots, k-1\}$, man e_i^2 and woman f_i^2 for $i \in \{2, \ldots, k\}$, and woman w^c for every clause c of ϕ. In addition, it contains the man-heavy edges (m_i, w_{i-1}) and (e_i^2, f_i^2) for $i \in \{2, \ldots, k\}$ and (m_i, f_{i-1}^2) for $i \in \{3, \ldots, k\}$, the balanced edges (m_1, w^c) for every clause c (tine edges), (e_i^1, w_i) for $i \in \{1, \ldots, k-1\}$ and (m_i, f_i^2) for $i \in \{2, \ldots, k\}$, and the woman-heavy edges (m_i, w_i) for $i \in \{1, \ldots, k\}$, and (e_i^1, f_i^1) and (e_i^1, w_{i+1}) for $i \in \{1, \ldots, k-1\}$.

Finally, instance $\mathcal{I}$ has a clause-accumulator connector (or *CA-connector*) for every clause c of ϕ consisting of the woman-heavy edge (m^c, w^c) between the man m^c (from the clause gadget corresponding to clause c) and woman w^c (from the accumulator); see Figure 4b. Notice that the above construction has more women than men. To restore balance, we pad the instance with extra (isolated) men that neither value nor are valued by any other agent. This completes the construction of the reduced instance.

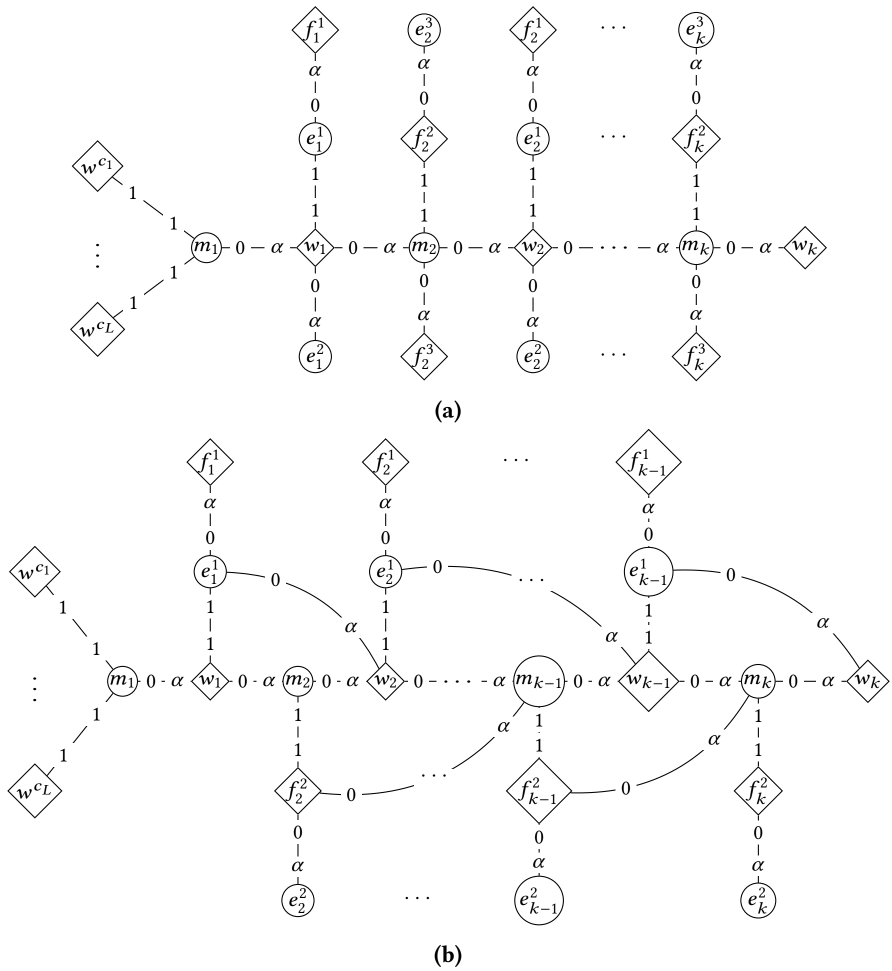

Fig. 6. *The accumulator for the cases (a) $\alpha \geq 2$ and (b) $\alpha \in (3/2, 2)$.*

5.2 Gadget properties

We will now prove several important properties (Claims 9-12) of our construction. The proofs appear in the full version [7].

CLAIM 9. *For every variable x, a stable fractional matching μ satisfies at least one of the following:*
(1) $\mu(m_1^x, w_1^{\overline{x}}) + \mu(m_1^x, w_2^{\overline{x}}) + \mu(m_1^x, f_1^x) = 1$ and $\mu(m_2^x, w_1^{\overline{x}}) + \mu(m_2^x, w_2^{\overline{x}}) + \mu(m_2^x, f_2^x) = 1$.
(2) $\mu(m_1^x, w_1^{\overline{x}}) + \mu(m_2^x, w_1^{\overline{x}}) + \mu(e_3^x, w_1^{\overline{x}}) = 1$ and $\mu(m_1^x, w_2^{\overline{x}}) + \mu(m_2^x, w_2^{\overline{x}}) + \mu(e_3^x, w_2^{\overline{x}}) = 1$.

We remark that the two conditions in the statement of Claim 9 affect the weight of the input edges of the VC-connectors that are attached to the variable gadget in any stable fractional matching. In particular, condition (1) implies that the weight assigned to the input edges of the VC-connectors that correspond to the two appearances of the positive literal x in clauses must be 0. To see why, observe that these input edges are incident to nodes m_1^x and m_2^x, and the total weight of all edges incident to each of these nodes cannot exceed 1. Condition (2) has a similar implication for the edges associated with the negative literal $\overline{x}$.

Claim 10. *Any stable fractional matching that assigns a weight of* 0 *to the input edge of a VC-connector must assign a weight of* 0 *to its output edge(s) as well.*

Claim 11. *Any stable fractional matching that assigns a weight of* 0 *to all output edges of the VC-connectors of clause* c *must assign a weight of* 0 *to the CA-connector of clause* c *as well.*

Claim 12. *Any stable fractional matching that assigns a weight of* 0 *to some CA-connector must assign a total weight of* 1 *to the tine edges and a weight of* 1 *to every balanced edge of the accumulator.*

5.3 Proof of inapproximability

Lemma 3. *If formula* ϕ *is not satisfiable, then any stable fractional matching of* $\mathcal{I}$ *has welfare at most* $80\alpha N + 4(k-1)$.

Proof. We will first show that if ϕ is not satisfiable, then any stable fractional matching of $\mathcal{I}$ assigns weight 0 to some CA-connector. For the sake of contradiction, consider a stable fractional matching that assigns a strictly positive weight to all CA-connectors. We will construct a truth assignment for the formula ϕ (contradicting the assumption of the lemma) by repeating the following process for every clause c of ϕ: Let ℓ be a literal that appears in c such that the output edge(s) of the VC-connector, that corresponds to the appearance of ℓ in c, have strictly positive total weight. By Claim 11, such a literal must exist. We set ℓ to 1 (true). For every variable that has not received a value in this way, we arbitrarily set it to 1.

The above assignment satisfies all the clauses. To show that it is also valid, we need to argue that there is no variable x such that both literals x and $\overline{x}$ have been set to 1. Assume, to the contrary, that literal x is set to 1 due to its appearance in a clause c_1, and literal $\overline{x}$ is set to 1 due to its appearance in a different clause c_2. Thus, in the above assignment, the output edge(s) of the VC-connector between the literal x and the clause c_1, as well as the VC-connector between the literal $\overline{x}$ and the clause c_2 have strictly positive (total) weight. By Claim 10, the input edges of both VC-connectors also have strictly positive weight. Let $i_1, i_2 \in \{1, 2\}$ be such that the i_1-th appearance of x is in the clause c_1 and the i_2-th appearance of $\overline{x}$ is in the clause c_2. Therefore, the said input edges are incident to the nodes $m^x_{i_1}$ and $w^{\overline{x}}_{i_2}$. Using Claim 9, we get that the total weight on the edges incident to one of $m^x_{i_1}$ or $w^{\overline{x}}_{i_2}$ exceeds 1, contradicting feasibility. Thus, the above assignment must be valid, which, in turn, implies that any stable fractional matching assigns weight 0 to some CA-connector.

By Claim 12, the contribution of the accumulator to the welfare is exactly $4k-2$ (2 from the tine edges plus 2 from each balanced edge). Let us now consider the contribution of the edges that do not belong to the accumulator. This comprises of

- a total value of 20 for the ten balanced edges of each of the N variable gadgets,
- a total value of α (respectively, $2+2\alpha$) for the edges of each of the $2N$ VC-connectors corresponding to a positive literal when $\alpha \geq 2$ (respectively, $\alpha \in (3/2, 2)$),
- a total value of $2+2\alpha$ (respectively, $6+6\alpha$) for the edges of each of the $2N$ VC-connectors corresponding to a negative literal when $\alpha \geq 2$ (respectively, $\alpha \in (3/2, 2)$),
- a total value of $18+\alpha$ for the nine balanced edges of each of the $4N/3$ clause gadgets and their corresponding CA-connectors.

It can be easily seen that $80\alpha N - 2$ is a (loose) upper bound on the total value from these edges. □

Lemma 4. *If* ϕ *is satisfiable, then there exists a stable fractional matching of* $\mathcal{I}$ *with welfare at least* $4(k-1)(\alpha - 1/2)$.

Proof. Starting from a satisfying assignment for ϕ, we will construct a stable fractional matching μ in which the welfare of the accumulator gadget is at least $4(k-1)(\alpha-1/2)$.

Variable gadgets. For the edges of the variable gadget of the variable x, μ is defined as:

- If x is true, then $\mu(m_1^x, w_1^{\overline{x}}) = \mu(e_3^x, w_1^{\overline{x}}) = \mu(e_3^x, w_2^{\overline{x}}) = \mu(m_2^x, w_2^{\overline{x}}) = 1/2$, $\mu(e_1^x, f_1^x) = \mu(e_2^x, f_2^x) = 1$, and the remaining edges have weight 0.
- If x is false, then $\mu(e_3^x, w_1^{\overline{x}}) = \mu(e_3^x, w_2^{\overline{x}}) = 1/2$, $\mu(m_1^x, f_1^x) = \mu(m_2^x, f_2^x) = 1$, and the remaining edges have weight 0.

Clause gadgets and CA-connectors. For each clause, select one of the true literals (tie-break arbitrarily) and call it *active*. Note that each clause has an active literal in a satisfying assignment. Consider the clause c, and let ℓ_i be its active literal for some $i \in \{1, 2, 3\}$. Also, let $i_1, i_2 \in \{1, 2, 3\} \setminus \{i\}$ denote the other two indices. Set $\mu(e_1^c, w_{i_1}^c) = \mu(e_2^c, w_{i_2}^c) = 1$, and set the weight of the remaining balanced edges to 0. Assign a weight of 1 to the CA-connector, i.e., $\mu(m^c, w^c) = 1$.

VC-connectors. For every non-active VC-connector, set the weight of its balanced edges (if any) to 1 and the weight of the remaining edges to 0. For every active VC-connector corresponding to the i-th appearance of the positive literal x as the j-th literal of clause c ($i \in \{1, 2\}, j \in \{1, 2, 3\}$), the weights of its edges are as follows:

- When $\alpha \geq 2$, we set $\mu(m_i^x, w_j^c) = 1/2$.
- When $\alpha \in (3/2, 2)$, we set $\mu(m_i^x, w^{x,c}) = 1/2$, $\mu(m^{x,c}, w^{x,c}) = 1 - \alpha/2$, and $\mu(m^{x,c}, w_j^c) = 1/\alpha$.

For every active VC-connector corresponding to the i-th appearance of the negative literal x as the j-th literal of clause c ($i \in \{1, 2\}, j \in \{1, 2, 3\}$), the weights of its edges are as follows:

- When $\alpha \geq 2$, we set $\mu(m^{\overline{x},c}, w_i^{\overline{x}}) = \mu(m^{\overline{x},c}, w_j^c) = 1/2$ and $\mu(m^{\overline{x},c}, w^{\overline{x},c}) = 0$.
- When $\alpha \in (3/2, 2)$, we set $\mu(m_1^{\overline{x},c}, w_i^{\overline{x}}) = 1/2$, $\mu(m_1^{\overline{x},c}, w_1^{\overline{x},c}) = 1 - \alpha/2$, $\mu(m_1^{\overline{x},c}, w_2^{\overline{x},c}) = (\alpha - 1)/2$, $\mu(m_2^{\overline{x},c}, w_2^{\overline{x},c}) = 1 - (\alpha^2 - \alpha)/2$, $\mu(m_2^{\overline{x},c}, w_j^c) = 2/\alpha - 1$, $\mu(m_3^{\overline{x},c}, w_1^{\overline{x},c}) = 1/\alpha$, $\mu(m_3^{\overline{x},c}, w_2^{\overline{x},c}) = \mu(m_3^{\overline{x},c}, w_3^{\overline{x},c}) = 0$, $\mu(m_3^{\overline{x},c}, w_j^c) = 1 - 1/\alpha$.

Accumulator. We set $\mu(m_1, w^c) = 0$ for every tine edge (m_1, w^c) of the accumulator. Furthermore:

- When $\alpha \geq 2$, we set $\mu(m_i, w_i) = 1/\alpha$ for all $i \in \{1, \ldots, k\}$, $\mu(e_i^2, w_i) = 1 - 2/\alpha$, $\mu(m_{i+1}, w_i) = 1/\alpha$, $\mu(e_i^1, f_i^1) = 1$, $\mu(e_i^1, w_i) = 0$ for all $i \in \{1, \ldots, k-1\}$, $\mu(m_i, f_i^2) = 0$, $\mu(m_i, f_i^3) = 1 - 2/\alpha$, and $\mu(e_i^3, f_i^2) = 1$ for all $i \in \{2, \ldots, k\}$. Among these, any edge with a positive weight is either man- or woman-heavy, and hence, its contribution to the social welfare is α times its weight. It can be verified that the total contribution is $4(k-1)(\alpha - 1/2) + 1$.
- When $\alpha \in (3/2, 2)$, we set $\mu(m_1, w_1) = 1/\alpha$, $\mu(m_2, w_2) = \alpha + 1/\alpha - 2$, $\mu(m_i, w_i) = 1 - 1/\alpha$ for all $i \in \{3, \ldots, k\}$, $\mu(m_{i+1}, w_i) = 1 - 1/\alpha$ for all $i \in \{1, \ldots, k-1\}$, $\mu(e_i^1, w_i) = 0$ for all $i \in \{1, \ldots, k-1\}$, $\mu(m_2, f_2^2) = 2 - \alpha$, $\mu(m_i, f_i^2) = 0$ for all $i \in \{3, \ldots, k\}$, $\mu(e_1^1, f_1^1) = \alpha - 1$, $\mu(e_2^2, f_2^2) = \alpha - 2/\alpha$, $\mu(e_k^2, f_k^2) = 1$, $\mu(e_i^1, f_i^1) = 2 - 2/\alpha$ for all $i \in \{2, \ldots, k-1\}$, $\mu(e_i^2, f_i^2) = 2 - 2/\alpha$ for all $i \in \{3, \ldots, k-1\}$, $\mu(e_1^1, w_2) = 2 - \alpha$, $\mu(e_i^1, w_{i+1}) = 2/\alpha - 1$ for all $i \in \{2, \ldots, k-1\}$, and $\mu(m_{i+1}, f_i^2) = 2/\alpha - 1$ for all $i \in \{2, \ldots, k-1\}$. Except for the balanced edge (m_2, f_2^2), every edge with a positive weight among the ones listed above is either man- or woman-heavy, and hence, its contribution to the social welfare is α times its weight. It can be verified that the total contribution in this case is $4(k-1)(\alpha - 1/2) + 2\alpha^2 - 7\alpha + 7$.

In each case, the accumulator contributes at least $4(k-1)(\alpha - 1/2)$ to the social welfare, as desired.

The feasibility of μ can be verified by inspection. To see why μ is stable, note that we only need to check for the balanced edges, as the man- or woman-heavy edges and the remaining pairs do not impose any constraints on stability. For the balanced edges, stability is established by the following series of observations (we will use the term 'stabilized by' to denote that an agent's utility is at least 1): The variable gadget for the variable x (Figure 4a) is stabilized by the agents f_1^x, f_2^x, e_3^x along with m_1^x, m_2^x (if x is true) or $w_1^{\overline{x}}, w_2^{\overline{x}}$ (if x is false). The clause gadget for clause c (Figure 4b) with active index i (and non-active indices i_1 and i_2) is stabilized by the agents e_1^c, e_2^c, w_i^c, $w_{i_1}^c$, $w_{i_2}^c$; in particular, the edge (m^c, w_i^c) is stabilized by w_i^c because an active literal triggers the woman-heavy edge in the VC-connector. A VC connector is stabilized by $w^{x,c}$ (Figure 5b), $m^{\overline{x},c}$ (Figure 5c), or $m_1^{\overline{x},c}$, $w_2^{\overline{x},c}$, and $m_3^{\overline{x},c}$ (Figure 5d). Finally, the tine edges in the accumulator (Figure 6) are stabilized by $w^{c_1}, \ldots, w^{c_L}$ (because we trigger the CA-connector), and the remaining balanced edges are stabilized by w_i's and m_i's except for m_1. Overall, μ is a feasible stable fractional matching. □

We are ready to prove Theorem 7. If $\alpha < N^{1+1/\delta}$, we use our construction with any k satisfying $k - 1 \geq \frac{20\alpha N(\alpha - 1/2 - \delta)}{\delta}$. It is easy to verify that the reduction is polynomial-time. Furthermore, from Lemma 3, we know that the welfare of μ when ϕ is not satisfiable is at most

$$80\alpha N + 4(k-1) \leq \frac{4(k-1)\delta}{\alpha - 1/2 - \delta} + 4(k-1) = \frac{4(k-1)(\alpha - 1/2)}{\alpha - 1/2 - \delta}.$$

This number is at least $\alpha - 1/2 - \delta$ times smaller than the welfare of μ when ϕ is satisfiable (Lemma 4). This establishes the inapproximability bound in part (i) of Theorem 7.

If $\alpha \geq N^{1+1/\delta}$, we use our construction with $k = N^{1+1/\delta}$. Once again, the reduction is polynomial-time, and the instance $\mathcal{I}$ has $n = \Theta(N^{1+1/\delta})$ men and women. Observe that $\alpha = \Omega(n)$, $k = \Theta(n)$, and $N = O(n^\delta)$. Hence, the welfare of μ when ϕ is not satisfiable is at most

$$80\alpha N + 4(k-1) \leq 80\alpha N + 4N^{1+1/\delta} \leq 84\alpha N = O(\alpha n^\delta).$$

On the other hand, the welfare of μ when ϕ is satisfiable is at least $4(k-1)(\alpha - 1/2)$, i.e., $\Omega(\alpha n)$. This establishes the bound in part (ii), and with it, completes the proof of Theorem 7.

ACKNOWLEDGMENTS

We are grateful to Elliot Anshelevich for bringing the work of Deligkas et al. [10] to our attention, to Argyrios Deligkas for sharing with us the full version of their paper [10], and to Haris Aziz for pointing us to the work of Manjunath [27]. Aris Filos-Ratsikas is supported by the Swiss National Science Foundation under contract No. 200021_165522, the ERC Advanced Grant 321171 (ALGAME). Ioannis Caragiannis and Aris Filos-Ratsikas acknowledge partial support by COST Action CA15210.

REFERENCES

[1] Atila Abdulkadiroğlu, Parag A Pathak, and Alvin E Roth. 2005. The New York City High School Match. *American Economic Review* (2005), 364–367.

[2] Elliot Anshelevich, Onkar Bhardwaj, and Martin Hoefer. 2017. Stable Matching with Network Externalities. *Algorithmica* 78 (2017), 1067–1106.

[3] Elliot Anshelevich and Sanmay Das. 2010. Matching, Cardinal Utility, and Social Welfare. *ACM SIGecom Exchanges* 9, 1 (2010), 4.

[4] Elliot Anshelevich, Sanmay Das, and Yonatan Naamad. 2013. Anarchy, Stability, and Utopia: Creating Better Matchings. *Autonomous Agents and Multi-Agent Systems* 26, 1 (2013), 120–140.

[5] Haris Aziz and Bettina Klaus. 2019. Random Matching under Priorities: Stability and No Envy Concepts. *Social Choice and Welfare* (2019), 1–47.

[6] Siddharth Barman. 2018. Approximating Nash Equilibria and Dense Subgraphs via an Approximate Version of Carathéodory's Theorem. *SIAM J. Comput.* 47, 3 (2018), 960–981.

[7] Ioannis Caragiannis, Aris Filos-Ratsikas, Panagiotis Kanellopoulos, and Rohit Vaish. 2019. Stable Fractional Matchings. *arXiv preprint arXiv:1902.06698* (2019).

[8] Xujin Chen, Guoli Ding, Xiaodong Hu, and Wenan Zang. 2012. The Maximum-Weight Stable Matching Problem: Duality and Efficiency. *SIAM Journal on Discrete Mathematics* 26, 3 (2012), 1346–1360.

[9] Ágnes Cseh, Robert W Irving, and David F Manlove. 2019. The Stable Roommates Problem with Short Lists. *Theory of Computing Systems* 63, 1 (2019), 128–149.

[10] Argyrios Deligkas, George B Mertzios, and Paul G Spirakis. 2017. The Computational Complexity of Weighted Greedy Matching. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, 317–325.

[11] Battal Doğan and Kemal Yıldız. 2016. Efficiency and Stability of Probabilistic Assignments in Marriage Problems. *Games and Economic Behavior* 95 (2016), 47–58.
[12] Federico Echenique and Alfred Galichon. 2017. Ordinal and Cardinal Solution Concepts for Two-Sided Matching. *Games and Economic Behavior* 101 (2017), 63–77.
[13] Federico Echenique, Alistair J Wilson, and Leeat Yariv. 2016. Clearinghouses for Two-Sided Matching: An Experimental Study. *Quantitative Economics* 7, 2 (2016), 449–482.
[14] Yuval Emek, Tobias Langner, and Roger Wattenhofer. 2015. The Price of Matching with Metric Preferences. In *Proceedings of the 23rd Annual European Symposium on Algorithms (ESA)*, Vol. 9294. Springer, 459–470.
[15] Tomás Feder. 1992. A New Fixed Point Approach for Stable Networks and Stable Marriages. *J. Comput. System Sci.* 45, 2 (1992), 233–284.
[16] David Gale and Lloyd S Shapley. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15.
[17] Dan Gusfield. 1987. Three Fast Algorithms for Four Problems in Stable Marriage. *SIAM J. Comput.* 16, 1 (1987), 111–128.
[18] Dan Gusfield and Robert W Irving. 1989. *The Stable Marriage Problem: Structure and Algorithms.* MIT Press.
[19] John William Hatfield and Scott Duke Kominers. 2015. Multilateral Matching. *Journal of Economic Theory* 156 (2015), 175–206.
[20] Robert W Irving. 1994. Stable Marriage and Indifference. *Discrete Applied Mathematics* 48, 3 (1994), 261–272.
[21] Robert W Irving, Paul Leather, and Dan Gusfield. 1987. An Efficient Algorithm for the "Optimal" Stable Marriage. *J. ACM* 34, 3 (1987), 532–543.
[22] Robert W Irving, David F Manlove, and Gregg O'Malley. 2009. Stable Marriage with Ties and Bounded Length Preference Lists. *Journal of Discrete Algorithms* 7, 2 (2009), 213–219.
[23] Kazuo Iwama and Shuichi Miyazaki. 2008. A Survey of the Stable Marriage Problem and its Variants. In *Proceedings of the International Conference on Informatics Education and Research for Knowledge-Circulating Society (ICKS).* IEEE, 131–136.
[24] Kazuo Iwama, Shuichi Miyazaki, Yasufumi Morita, and David Manlove. 1999. Stable Marriage with Incomplete Lists and Ties. In *Proceedings of the 26th International Colloquium on Automata, Languages, and Programming (ICALP).* Springer, 443–452.
[25] Donald Ervin Knuth. 1997. *Stable Marriage and its Relation to Other Combinatorial Problems: An Introduction to the Mathematical Analysis of Algorithms.* Vol. 10. American Mathematical Soc.
[26] Tung Mai and Vijay V Vazirani. 2018. A Natural Generalization of Stable Matching Solved via New Insights into Ideal Cuts. *arXiv preprint arXiv:1802.06621* (2018).
[27] Vikram Manjunath. 2013. Stability and the Core of Probabilistic Marriage Problems. (2013).
[28] David F Manlove. 2013. *Algorithmics of Matching under Preferences.* Vol. 2. World Scientific.
[29] David F Manlove, Robert W Irving, Kazuo Iwama, Shuichi Miyazaki, and Yasufumi Morita. 2002. Hard Variants of Stable Marriage. *Theoretical Computer Science* 276, 1-2 (2002), 261–279.
[30] Maria Silvia Pini, Francesca Rossi, K Brent Venable, and Toby Walsh. 2013. Stability, Optimality and Manipulation in Matching Problems with Weighted Preferences. *Algorithms* 6, 4 (2013), 782–804.
[31] Alvin E Roth and Elliott Peranson. 1999. The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design. *American Economic Review* 89, 4 (1999), 748–780.
[32] Alvin E Roth, Uriel G Rothblum, and John H Vande Vate. 1993. Stable Matchings, Optimal Assignments, and Linear Programming. *Mathematics of Operations Research* 18, 4 (1993), 803–828.
[33] Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. 2004. Kidney Exchange. *The Quarterly Journal of Economics* 119, 2 (2004), 457–488.
[34] Alvin E Roth and Marilda A Oliveira Sotomayor. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis.* Cambridge University Press.
[35] Uriel G Rothblum. 1992. Characterization of Stable Matchings as Extreme Points of a Polytope. *Mathematical Programming* 54, 1-3 (1992), 57–67.
[36] Chung-Piaw Teo and Jay Sethuraman. 1998. The Geometry of Fractional Stable Matchings and its Applications. *Mathematics of Operations Research* 23, 4 (1998), 874–891.
[37] John H Vande Vate. 1989. Linear Programming Brings Marital Bliss. *Operations Research Letters* 8, 3 (1989), 147–153.
[38] Rakesh V Vohra. 2012. Stable Matchings and Linear Programming. *Current Science* 105, 9 (2012), 1051–1055.
[39] Ryo Yoshinaka. 2005. Higher-order Matching in the Linear Lambda Calculus in the Absence of Constants is NP-complete. In *Proceedings of the 16th International Conference on Term Rewriting and Applications (RTA).* Springer-Verlag, 235–249.