

האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

**A STABLE MARRIAGE REQUIRES
COMMUNICATION**

By

**YANNAI A. GONCZAROWSKI
and NOAM NISAN**

Discussion Paper # 667 (June 2014)

מרכז לחקר הרציונליות

**CENTER FOR THE STUDY
OF RATIONALITY**

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: <http://www.ratio.huji.ac.il/>

A Stable Marriage Requires Communication

Yannai A. Gonczarowski* Noam Nisan†

June 1, 2014

Abstract

The Gale-Shapely algorithm for the Stable Marriage Problem is known to take $\Theta(n^2)$ steps to find a stable marriage *in the worst case*, but only $\Theta(n \log n)$ steps *in the average case* (with n women and n men). In 1976, Knuth asked whether the worst-case running time can be improved in a model of computation that does not require sequential access to the whole input. A partial negative answer was given by Ng and Hirschberg, who showed that $\Theta(n^2)$ queries are required in a model that allows certain natural random-access queries to the participants' preferences.

Using a reduction to the communication complexity of the disjointness problem, we prove a significantly more general — albeit slightly weaker — result, showing that $\Omega(n^2)$ Boolean queries of any type are required. Our lower bound generalizes to (A) randomized algorithms, (B) even just verifying the stability of a proposed marriage, (C) even allowing arbitrary separate preprocessing of the women's preferences and of the men's preferences, and (D) several variants of the basic problem, such as whether a given pair is married in every/some stable marriage.

1 Introduction

In the classic Stable Marriage Problem (Gale and Shapley, 1962), there are n women and n men; each woman has a full preference order over the men and each man has a full preference order over the women. The challenge is to find a *stable marriage*: a 1:1 mapping between women and men that is stable in the sense of having no *blocking pair*: a woman and man who both prefer each other over their current spouse in the marriage. Gale and Shapley (1962) proved that such a stable marriage exists, by providing an algorithm for finding one. Their algorithm takes $\Theta(n^2)$ steps in the worst case (Gale and Shapley, 1962), but only $\Theta(n \log n)$ steps in the average case, over independently and uniformly chosen preferences (Wilson, 1972).

In 1976, Knuth asked whether this quadratic worst-case running time can be improved upon. An easier related question was put forward by Gusfield in 1987, who asked whether even *verifying* the stability of a proposed marriage can be done any faster. As the input size here is quadratic, these questions only make sense in models that do not require sequentially reading the whole input, but rather provide some kind of random access to the preferences of the participants.

*The Hebrew University of Jerusalem (Institute of Mathematics, School of Computer Science & Engineering and Center for the Study of Rationality) and Microsoft Research.

†The Hebrew University of Jerusalem (School of Computer Science & Engineering and Center for the Study of Rationality) and Microsoft Research.

A partial answer to both questions was given by Ng and Hirschberg (1990), who considered a model that allows two types of unit-cost queries to the preferences of the participants: “what is woman w ’s ranking of man m ?” (and, dually, “what is man m ’s ranking of woman w ?”) and “which man does woman w rank at place k ?” (and, dually, “which woman does man m rank at place k ?”). In this model, they prove a tight $\Theta(n^2)$ lower bound on the number of queries that any algorithm that solves the stable marriage problem, or even verifies whether a given marriage is stable, must make in the worst case.

These results of Ng and Hirschberg (1990) still leave two questions open. The first is whether some more powerful model may allow for faster algorithms. While most “natural” algorithms for stable marriage do fit into this model, there may be others that do not; indeed, there exist problems for which “unnatural” operations, such as various types of hashing or arithmetic operations, do give algorithmic speedups. The second question concerns randomized algorithms: can they do better than deterministic ones here? This question is especially fitting for this problem as the *expected* running time is known to be small.¹

We give a negative answer to both hopes, using a reduction to the well-known lower bounds for the disjointness problem (Kalyanasundaram and Schintger, 1992; Razborov, 1992) in Yao’s model of two-party communication complexity (see Kushilevitz and Nisan, 1997, for a survey). Consider a scenario in which Alice holds the preferences of the n women and Bob holds the preferences of the n men. In Section 4, we show that Alice and Bob must exchange $\Omega(n^2)$ bits in the worst case even for the verification problem and even if they are allowed to use randomized algorithms with bounded error; in Section 5, we show that the same lower bound holds also for finding a stable marriage. These results immediately imply the same lower bounds for any type of Boolean queries in the original model.

Theorem 1.1 (Informal version of Theorems 3.1 and 3.3). *For every type of Boolean queries to the women’s preferences and to the men’s preferences, every randomized (or deterministic) algorithm for producing a stable marriage must make $\Omega(n^2)$ queries in the worst case. The same lower bound applies also to verifying the stability of a proposed marriage given as an additional input.*

The lower bound holds regardless of which of the stable marriages is produced by the algorithm; it also holds even if we allow for arbitrary preprocessing of all the women’s preferences (separately from the men’s), and dually of all of the men’s preferences (separately from the women’s). In Section 6, we also prove the same lower bound for various variants of the problem, such as the decision problem of whether a stable marriage exists whence some given woman and man are married to each other.

The lower bound for verification complexity is tight, and indeed there exists a simple deterministic algorithm for verifying the stability of a proposed marriage, which requires $O(n^2)$ queries even in the weak *comparison model* that allows only for queries of the form “does woman w prefer man m_1 over man m_2 ?” and, dually, “does man m prefer woman w_1 over woman w_2 ?”.² We do not know whether the lower bound is tight also for

¹In particular, this would be the case if the expected running time could be made small for *any* distribution on preferences, rather than just the uniform one.

²By simple batching, this verification algorithm can be converted into one that uses only $O(n^2/\log n)$ queries, each of which returns an answer of length $\log n$ bits (with each query still regarding the preferences of only a single participant). This highlights the fact that the lower bounds of Ng and Hirschberg (1990) crucially depend on the exact type of queries allowed in their model.

finding a stable marriage. Gale and Shapley’s algorithm uses $O(n^2)$ *queries* in the worst case, but $O(n^2)$ of these queries require each an answer of length $\log n$ bits, and thus the algorithm requires a total of $O(n^2 \log n)$ *Boolean queries*, or bits of communication. We do not know whether $O(n^2)$ *Boolean queries* suffice for any algorithm. While the gap between Gale and Shapley’s algorithm and our lower bound is small, we believe that it is interesting, as the number of queries performed by the algorithm is exactly linear in the input encoding length; an even slightly sublinear algorithm would therefore be interesting.³ We indeed do not have any $o(n^2 \log n)$ algorithm, even randomized and even in the strong two-party communication model, nor do we have any improved $\omega(n^2)$ lower bound, even for deterministic algorithms and even in the simple comparison model.

Open Problem 1. Consider the Comparison model for stable marriage that only allows for queries of the form “does man m prefer woman w_1 over woman w_2 ?” and, dually, “does woman w prefer man m_1 over man m_2 ?”. How many such queries are required, in the worst case, to find a stable marriage?

2 Model and Preliminaries

2.1 The Stable Marriage Problem

2.1.1 Full Preference Lists

For ease of presentation, we consider a simplified version of the model of Gale and Shapley (1962). Let W and M be disjoint finite sets, denoted *women* and *men*, respectively, s.t. $|W| = |M|$.

Definition 2.1 (Full Preferences).

1. A *full preference list* over M is a total ordering of M .
2. A *profile of full preference lists* for W over M is a specification of a full preference list over M for each woman $w \in W$. We denote the set of all profiles of full preference lists for W over M by $\mathcal{F}(W, M)$.
3. Given a profile P_W of full preference lists for W over M , a woman $w \in W$ is said to *prefer a man* $m \in M$ *over a man* $m' \in M$, denoted by $m \succ_w m'$, if m precedes m' on the preference list of w .

We define full preference lists over W and profiles of full preference lists for M over W analogously.

Definition 2.2 (Perfect Marriage). A *perfect marriage* between W and M is a one-to-one mapping between W and M .

Definition 2.3 (Stability). Let P_W and P_M be profiles of full preference lists for W over M and for M over W , respectively. A perfect marriage μ is said to be *unstable* (w.r.t. P_W and P_M) if there exist a woman $w \in W$ and a man $m \in M$, each preferring the other over the partner married to them in μ . If μ is not unstable, then it is said to be *stable*.

³Note that, as shown in Appendix B, the *nondeterministic* communication complexity is $\Theta(n^2)$, so proving higher lower bounds for the deterministic or randomized case may be challenging.

2.1.2 Arbitrary Preference Lists

While our main results are phrased in terms of full preference lists and perfect marriages, some additional and intermediate results in Sections 5 and 6 (as well as in Appendices A and B) deal with a more extended model, which allows for preferences to specify “black-lists” (i.e. declare some potential partners as unacceptable) and for marriages to specify that some participants remain single. (This model is nonetheless also a simplified version of that of Gale and Shapley (1962).) A (not necessarily full) *preference list* over M is a totally-ordered subset of M . We once again interpret a preference list as a ranking, from best to worst, of acceptable partners; we interpret participants absent from a preference list as declared unacceptable, even at the cost of remaining single. Analogously, a *profile of preference lists* for W over M is a specification of a preference list over M for each woman $w \in W$; we denote the set of all profiles of preference lists for W over M by $\mathcal{P}(W, M) \supset \mathcal{F}(W, M)$. In this more extended model, a woman w is said to *prefer a man m over a man m'* not only when m precedes m' on the preference list of w , but also when m is on the preference list of w while m' is not. (We once again define preference lists and profiles of preference list for M over W analogously.)

A (not necessarily perfect) *marriage* between W and M is a one-to-one mapping between a subset of W and a subset of M . Given a marriage μ , we denote the set of married women (i.e. the subset of W on which μ is defined) by W_μ ; we analogously denote the set of married men by M_μ . A marriage μ is said to be *unstable* (w.r.t P_W and P_M) not only if some pair would rather deviate, but also if some participant $p \in W \cup M$ is married to someone not on the preference list of p .

We note that this model indeed generalizes the one from Section 2.1.1, in the sense that when the preference list of every participant contains all participants of the other side, then the definition of a stable marriage in this extended model (w.r.t. these preference lists) coincides with that of the simpler model (w.r.t. these preference lists when viewed as full preference lists). In particular, any marriage that is stable w.r.t. such preference lists prescribes for no participant to remain single.

2.1.3 Known Results

We now survey a few known results regarding the stable marriage problem, which we utilize throughout this paper. For the duration of this section, let $P_W \in \mathcal{P}(W, M)$ and $P_M \in \mathcal{P}(M, W)$ be profiles of preference lists for W over M and for M over W , respectively.

Theorem 2.4 (Gale and Shapley (1962)). *A stable marriage between W and M always exists. Moreover, there exists an M -optimal stable marriage, i.e. a stable marriage that is weakly preferred by each man over every other stable marriage.*

Theorem 2.5 (McVitie and Wilson (1971)). *The M -optimal stable marriage is also the W -worst stable marriage, i.e. every other stable marriage is weakly preferred over it by each woman.*

Corollary 2.6 (W -worst & M -worst \Rightarrow unique). *If a stable marriage is both the W -worst stable marriage and the M -worst stable marriage, then it is the unique stable marriage.*

Theorem 2.7 (Rural Hospitals Theorem (Roth, 1986)). *W_μ (resp. M_μ) is the same for every stable marriage μ .*

2.2 Communication Complexity

We work in Yao’s (1979) model of two-party communication complexity (see Kushilevitz and Nisan, 1997, for a survey). Consider a scenario where Alice holds a value x , Bob holds a value y , and Alice and Bob wish to perform some computation that depends on both x and y ; such a computation generally requires the exchange of some information between Alice and Bob. The *communication complexity* of a given protocol (i.e. distributed algorithm) for such a computation is the number of bits that Alice and Bob exchange under this protocol in the worst case (i.e. for the worst x, y); the *communication complexity* of the computation that Alice and Bob wish to perform is the lowest of the communication complexities of any protocol for this computation. Generalizing, we also consider *randomized* communication complexity, defined analogously using randomized protocols that for every given fixed input, produce a correct output with probability at least $2/3$.⁴

In our proofs, we make use of the known communication complexity of calculating the disjointness function.

Theorem 2.8 (Kalyanasundaram and Schintger (1992); see also Razborov (1992)). *Let $n \in \mathbb{N}$. The randomized (and deterministic) communication complexity of calculating $\bigvee_{i=1}^n x_i y_i$, where $\bar{x} = (x_i)_{i=1}^n$ is a sequence of n bits held by Alice, and $\bar{y} = (y_i)_{i=1}^n$ is a sequence of n bits held by Bob, is $\Theta(n)$.*

3 Main Results

All of our results provide lower bounds on various computations regarding the stable marriage problem. For the duration of this section, let $n \in \mathbb{N}$, let $W = \{w_1, \dots, w_n\}$ and $M = \{m_1, \dots, m_n\}$ be disjoint sets s.t. $|W| = |M| = n$, and let $P_W \in \mathcal{F}(W, M)$ and $P_M \in \mathcal{F}(M, W)$ be profiles of full preference lists for W over M and for M over W , respectively.

Theorem 3.1 (Complexity of Verification of Stability). *Let μ_{id} be the perfect marriage in which w_i is married to m_i for every i . For every type of Boolean queries that query the women’s preferences separately from the men’s preferences, every randomized (or deterministic) algorithm for determining whether μ_{id} is stable must make $\Omega(n^2)$ queries in the worst case.*

Remark 3.2. The lower bound given in Theorem 3.1 is tight. Indeed, exhausting over all pairs to naïvely check for the existence of a blocking pair requires $\Theta(n^2)$ Boolean queries in the worst case.

A proof of Theorem 3.1 is given in Section 4.

Theorem 3.3 (Complexity of Finding a Stable Marriage — Lower Bound). *For every type of Boolean queries that query the women’s preferences separately from the men’s preferences, every randomized (or deterministic) algorithm for producing a stable marriage must make $\Omega(n^2)$ queries in the worst case.*

A proof of Theorem 3.3 is given in Section 5. As noted in Section 1, the question of a tight lower bound for finding a stable marriage remains open — see Open Problem 1.

⁴The results of this paper hold verbatim even if the constant $2/3$ is replaced with any other fixed probability $1/2 < p \leq 1$.

Theorem 3.4 (Complexity of Determining the Marital Status of a Given Couple — Lower Bound).

1. For every type of Boolean queries that query the women's preferences separately from the men's preferences, every randomized (or deterministic) algorithm for determining whether w_1 and m_1 are married in **some** stable marriage between W and M must make $\Omega(n^2)$ queries in the worst case.
2. For every type of Boolean queries that query the women's preferences separately from the men's preferences, every randomized (or deterministic) algorithm for determining whether w_1 and m_1 are married in **every** stable marriage between W and M must make $\Omega(n^2)$ queries in the worst case.

A proof of Theorem 3.4 is given in Section 6. Once again, the question of tight lower bounds for both of these problems remains open.⁵

4 Verification of Stability

In this section, we prove Theorem 3.1 and introduce the general construction technique underlying all of the results of this paper. We continue to use the notation of Section 3 throughout this section as well. Theorem 3.1 directly follows from the following theorem regarding the communication complexity of verifying the stability of a given marriage.

Theorem 4.1 (Communication Complexity of Verification of Stability). *The randomized (and deterministic) communication complexity of determining whether μ_{id} is stable, where P_W is held by Alice and P_M is held by Bob, is $\Omega(n^2)$.*

Remark 4.2. Theorem 4.1 is phrased so that the marriage $\mu = \mu_{\text{id}}$ is known by both Alice and Bob before the protocol commences. Nonetheless, this theorem still holds if only one of them knows μ , as the straightforward way of encoding a marriage between W and M requires $O(n \log n)$ bits.

We prove Theorem 4.1 by embedding the problem of disjointness in that of verification of stability (Lemma 4.4), and then applying Theorem 2.8

Definition 4.3. We denote the set of pairs of distinct elements of $\{1, \dots, n\}$ by

$$[n]^{\bar{2}} \triangleq \{(i, j) \in \{1, \dots, n\}^2 \mid i \neq j\}.$$

We note that $|[n]^{\bar{2}}| = n \cdot (n - 1)$.

Lemma 4.4 (Disjointness \leftrightarrow Verification of Stability). *There exist functions $P_W : \{0, 1\}^{[n]^{\bar{2}}} \rightarrow \mathcal{F}(W, M)$ and $P_M : \{0, 1\}^{[n]^{\bar{2}}} \rightarrow \mathcal{F}(M, W)$ s.t. for every $\bar{x} = (x_j^i)_{(i,j) \in [n]^{\bar{2}}} \in \{0, 1\}^{[n]^{\bar{2}}}$ and $\bar{y} = (y_j^i)_{(i,j) \in [n]^{\bar{2}}} \in \{0, 1\}^{[n]^{\bar{2}}}$, the following are equivalent.*

- μ_{id} is stable w.r.t. $P_W(\bar{x})$ and $P_M(\bar{y})$
- $\bigvee_{(i,j) \in [n]^{\bar{2}}} x_j^i y_j^i = 0$.

⁵Gusfield (1987) gives a deterministic algorithm for enumerating all pairs that belong to at least one stable marriage in $O(n^2 \log n)$ Boolean queries; this yields a $O(n^2 \log n)$ upper bound for both problems.

Before proving Lemma 4.4, we first show how Theorem 4.1 follows from it.

Proof of Theorem 4.1. Every protocol for verification of stability of μ_{id} may be used to construct a protocol, with the same communication complexity, for disjointness as follows: Alice computes $P_W(\bar{x})$, Bob computes $P_M(\bar{y})$, and they use the given protocol to determine whether μ_{id} is stable; by Lemma 4.4, $\bigvee_{(i,j) \in [n]^2} x_j^i y_j^i = 0$ iff μ_{id} is stable. By Theorem 2.8, the communication complexity of the resulting protocol is $\Omega(n^2)$, and therefore so is that of the given protocol for verification of stability. \square

We conclude this section by proving Lemma 4.4.

Proof of Lemma 4.4. To define $P_W(\bar{x})$, for every i we define the preference list of w_i to consist of all m_j s.t. $x_j^i = 1$, in arbitrary order (say, sorted by j), followed by m_i , followed by all other men in arbitrary order. Similarly, to define $P_M(\bar{y})$, for every j we define the preference list of m_j to consist of all w_i s.t. $y_j^i = 1$, in arbitrary order (say, sorted by i), followed by w_j , followed by all other women arbitrary order.

μ_{id} is unstable w.r.t. $P_W(\bar{x})$ and $P_M(\bar{y}) \iff$ there exist $(i, j) \in [n]^2$ s.t. $m_j \succ_{w_i} m_i$ and $w_i \succ_{m_j} w_j \iff$ there exist $(i, j) \in [n]^2$ s.t. $x_j^i = 1$ and $y_j^i = 1 \iff \bigvee_{(i,j) \in [n]^2} x_j^i y_j^i \neq 0$. \square

Remark 4.5. A similar argument may be used to embed verification of stability back in disjointness.

5 Finding a Stable Marriage

In this section, we prove Theorem 3.3. This theorem follows directly from the following theorem regarding the communication complexity of finding a stable marriage.

Theorem 5.1 (Communication Complexity of Finding a Stable Marriage — Lower Bound). *Let $n \in \mathbb{N}$, let W and M be disjoint sets s.t. $|W| = |M| = n$, and let $P_W \in \mathcal{F}(W, M)$ and $P_M \in \mathcal{F}(M, W)$. The randomized (and deterministic) communication complexity of finding a stable marriage, where P_W is held by Alice and P_M is held by Bob, is $\Omega(n^2)$.*

Similarly to the derivation of Theorem 4.1, we prove Theorem 5.1 by embedding the problem of disjointness in that of finding a stable marriage, and then applying Theorem 2.8; we perform this embedding through the intermediate problem of finding a stable marriage w.r.t. arbitrary (i.e. not necessarily full) preference lists.

Lemma 5.2 (Disjointness \leftrightarrow Finding a Stable Marriage (Arbitrary Preferences)). *Let $n \in \mathbb{N}$, let $W = \{w_1, \dots, w_n\}$ and $M = \{m_1, \dots, m_n\}$ be disjoint sets s.t. $|W| = |M| = n$. Let μ_{id} be the perfect marriage in which w_i is married to m_i for every i . There exist functions $P_W : \{0, 1\}^{[n]^2} \rightarrow \mathcal{P}(W, M)$ and $P_M : \{0, 1\}^{[n]^2} \rightarrow \mathcal{P}(M, W)$ s.t. for every $\bar{x} = (x_j^i)_{(i,j) \in [n]^2} \in \{0, 1\}^{[n]^2}$ and $\bar{y} = (y_j^i)_{(i,j) \in [n]^2} \in \{0, 1\}^{[n]^2}$, both of the following hold.*

1. *If $\bigvee_{(i,j) \in [n]^2} x_j^i y_j^i = 0$, then μ_{id} is the unique stable marriage w.r.t. $P_W(\bar{x})$ and $P_M(\bar{y})$.*
2. *If $\bigvee_{(i,j) \in [n]^2} x_j^i y_j^i \neq 0$, then μ_{id} is unstable w.r.t. $P_W(\bar{x})$ and $P_M(\bar{y})$.*

Proof. To define $P_W(\bar{x})$, for every i we define the preference list of w_i to consist of all m_j s.t. $x_j^i = 1$, in arbitrary order (say, sorted by j), followed by m_i (with all other men absent). Similarly, to define $P_M(\bar{y})$, for every j we define the preference list of m_j to

consist of all w_i s.t. $y_j^i = 1$, in arbitrary order (say, sorted by i), followed by w_j (with all other women absent).

We first show that μ_{id} is stable w.r.t. $P_W(\bar{x})$ and $P_M(\bar{y})$ iff $\bigvee_{(i,j) \in [n]^2} x_j^i y_j^i = 0$. Indeed, similarly to the proof of Lemma 4.4 and since every participant is married by μ_{id} to someone on their preference list, we have: μ_{id} is unstable w.r.t. $P_W(\bar{x})$ and $P_M(\bar{y}) \iff$ there exist $(i, j) \in [n]^2$ s.t. $m_j \succ_{w_i} m_i$ and $w_i \succ_{m_j} w_j \iff$ there exist $(i, j) \in [n]^2$ s.t. $x_j^i = 1$ and $y_j^i = 1 \iff \bigvee_{(i,j) \in [n]^2} x_j^i y_j^i \neq 0$.

It remains to show that if μ_{id} is stable w.r.t. $P_W(\bar{x})$ and $P_M(\bar{y})$, then it is the unique stable marriage w.r.t. these profiles of preference lists. For the remainder of the proof assume, therefore, that μ_{id} is stable (w.r.t. $P_W(\bar{x})$ and $P_M(\bar{y})$); let μ be a stable marriage (w.r.t. these profiles of preference lists). As μ_{id} is stable and perfect, by Theorem 2.7 we have that since μ is stable, it is perfect as well. Therefore, each $p \in W \cup M$ is married by μ to someone on the preference list of p , and so p weakly prefers μ over μ_{id} , as in the latter p is married to the last person on the preference list of p . Thus, μ_{id} is both the W -worst stable marriage and the M -worst stable one, and so, by Corollary 2.6, μ_{id} is the unique stable marriage. \square

Definition 5.3 (Submarriage). Let W' and M' be disjoint sets. A marriage μ , between a subset W of W' and a subset M of M' , is said to be a *submarriage* of a marriage μ' between W' and M' , if for every $w \in W$ and $m \in M$, we have $\mu'(w) = m$ iff $\mu(w) = m$.

Lemma 5.4 (Finding a Stable Marriage (Arbitrary Preferences) \iff Finding a Stable Marriage (Full Preferences)). *Let $n \in \mathbb{N}$, and let W, W', M and M' be pairwise-disjoint sets, each of cardinality n . There exist functions $P_{W \cup W'} : \mathcal{P}(W, M) \rightarrow \mathcal{F}(W \cup W', M \cup M')$ and $P_{M \cup M'} : \mathcal{P}(M, W) \rightarrow \mathcal{F}(M \cup M', W \cup W')$ s.t. for every $P_W \in \mathcal{P}(W, M)$ and $P_M \in \mathcal{P}(M, W)$, and for every (possibly imperfect) marriage μ between W and M , the following are equivalent.*

- μ is stable w.r.t. P_W and P_M .
- μ is a submarriage of some marriage between $W \cup W'$ and $M \cup M'$ that is stable w.r.t. $P_{W \cup W'}(P_W)$ and $P_{M \cup M'}(P_M)$.

*Proof.*⁶ Denote $W = \{w_1, \dots, w_n\}$, $M = \{m_1, \dots, m_n\}$, $W' = \{w'_1, \dots, w'_n\}$, and $M' = \{m'_1, \dots, m'_n\}$.

To define $P_{W \cup W'}(P_W)$, for every i we define the preference list of w_i to consist of her preference list in P_W (in the same order), followed by m'_i , followed by all other men in arbitrary order; we define the preference list of w'_i to consist of m_i , followed by all other men in arbitrary order. Similarly, to define $P_{M \cup M'}(P_M)$, for every j we define the preference list of m_j to consist of his preference list in P_M (in the same order), followed by w'_j , followed by all other women in arbitrary order; we define the preference list of m'_j to consist of w_j , followed by all other women in arbitrary order.

It is straightforward to verify that the lemma holds w.r.t. these definitions of $P_{W \cup W'}$ and $P_{M \cup M'}$; the details are left to the reader. \square

Remark 5.5. It is straightforward to embed the problem of finding a stable marriage w.r.t. full preference lists in that of finding a stable marriage w.r.t. arbitrary preference lists, as the former is a special case of the latter.

⁶Our construction in this proof is essentially a one-to-one version of the many-to-many construction given in Corollary 31 of Gonczarowski and Friedgut (2013).

The proof of Theorem 5.1 given Lemmas 5.2 and 5.4 is similar to that of Theorem 4.1. Given a protocol for finding a stable marriage, a protocol for disjointness with the same communication complexity can be constructed as follows: Alice computes $P_A \triangleq P_{W \cup W'}(P_W(\bar{x}))$, Bob computes $P_B \triangleq P_{M \cup M'}(P_M(\bar{y}))$, and they use the given protocol to find a marriage μ that is stable w.r.t. P_A and P_B ; by Lemmas 5.2 and 5.4, we have that $\bigvee_{(i,j) \in [n]^2} x_j^i y_j^i = 0$ iff μ_{id} is a submarriage of μ .

6 Determining the Marital Status of a Given Couple or Participant

In this section, we prove Theorem 3.4. This theorem follows directly from the following theorem regarding the communication complexity of the relevant problems.

Theorem 6.1 (Communication Complexity of Determining the Marital Status of a Given Couple — Lower Bound). *Let $n \in \mathbb{N}$, let $W = \{w_1, \dots, w_n\}$ and $M = \{m_1, \dots, m_n\}$ be disjoint sets s.t. $|W| = |M| = n$, and let $P_W \in \mathcal{F}(W, M)$ and $P_M \in \mathcal{F}(M, W)$. The randomized (and deterministic) communication complexity of determining whether w_1 and m_1 are married in some (alternatively, in every) stable marriage, where P_W is held by Alice and P_M is held by Bob, is $\Omega(n^2)$.*

We prove Theorem 6.1 once again using Theorem 2.8, by embedding disjointness in both problems. We do so by embedding disjointness in an intermediate problem, and in turn embedding this intermediate problem in both problems; this intermediate problem is that of determining whether a given participant is single (i.e. not married to anyone) in some stable marriage, given profiles of arbitrary (i.e. not necessarily full) preference lists.⁷ We therefore conclude analogues of Theorems 3.4 and 6.1 for this problem as well. (See Corollary 6.3.)

Lemma 6.2 (Disjointness \leftrightarrow Is Participant Single?). *Let $n \in \mathbb{N}$, let W and M be disjoint sets s.t. $|W| = |M| = 2n$, and let $p \in W \cup M$. There exist functions $P_W : \{0, 1\}^{[n]^2} \rightarrow \mathcal{P}(W, M)$ and $P_M : \{0, 1\}^{[n]^2} \rightarrow \mathcal{P}(M, W)$ s.t. for every $\bar{x} = (x_j^i)_{(i,j) \in [n]^2} \in \{0, 1\}^{[n]^2}$ and $\bar{y} = (y_j^i)_{(i,j) \in [n]^2} \in \{0, 1\}^{[n]^2}$, the following are equivalent.*

- p is single in some stable marriage w.r.t. $P_W(\bar{x})$ and $P_M(\bar{y})$.
- $\bigvee_{(i,j) \in [n]^2} x_j^i y_j^i = 0$.

Proof. Assume w.l.o.g. that $p \in W$ and denote $w \triangleq p$. Denote $W = \{w_1, \dots, w_n, w, w'_2, w'_3, \dots, w'_n\}$ and $M = \{m_1, \dots, m_n, m'_1, \dots, m'_n\}$.

To define $P_W(\bar{x})$, for every i we define the preference list of w_i to consist of all m_j s.t. $x_j^i = 1$, in arbitrary order (say, sorted by j), followed by m'_i (with all other men absent). We define the preference list of w to consist of all m'_j , in arbitrary order (say, sorted by j), with all other men absent. We define the preference list of every w'_i to be empty (these women can be ignored, and are defined purely for aesthetic reasons — so that W and M be of equal cardinality). To define $P_M(\bar{y})$, for every j we define the preference list of m_j to consist of all w_i s.t. $y_j^i = 1$, in arbitrary order (say, sorted by i),

⁷By Theorem 2.7 (in conjunction with Theorem 2.4), this is equivalent to whether this participant is single in *every* stable marriage.

with all other women absent. For every j we define the preference list of m'_j to consist of w_j , followed by w (with all other women absent).

Let μ'_{id} be the marriage in which w_i is married to m'_i for every i , and in which all other participants are single. We first show that $\bigvee_{(i,j) \in [n]^2} x_j^i y_j^i = 0$ iff μ'_{id} is stable, and then show that μ'_{id} is stable iff $w = p$ is single in some stable marriage; we commence with the former.

We begin by noting that every participant that is married in μ'_{id} is married to someone on their preference list; therefore, μ'_{id} is stable iff no pair would rather deviate. Obviously, no w'_i would rather deviate with anyone. Furthermore, while w would rather deviate with any m'_j , these are all married to their top choices, and so none of them would deviate with w . Since for every i , the preference list of w_i consists of m'_i and of a subset of $\{m_j\}_{j \neq i}$, we therefore have that μ'_{id} is unstable iff there exists $(i, j) \in [n]^2$ s.t. both $m_j \succ_{w_i} m'_i$ and w_i is on the preference list of m_j . Similarly to the proofs of Lemmas 4.4 and 5.2, this holds precisely if there exists $(i, j) \in [n]^2$ s.t. $x_j^i = 1$ and $y_j^i = 1$, which holds iff $\bigvee_{(i,j) \in [n]^2} x_j^i y_j^i \neq 0$.

We complete the proof by showing that μ'_{id} is stable iff $w = p$ is single in some stable marriage. The first direction follows immediately from the fact that w is single in μ'_{id} . For the second direction, assume that there exists a stable marriage μ in which w is single. By stability of μ and since all men on the preference list of w have w on their preference list, all such men are married in μ and prefer their spouses over w . Therefore, for every j , we have that m'_j is married to w_j in μ . By stability of μ , every w'_i is single in μ . As μ and μ'_{id} coincide on all women, we have that $\mu = \mu'_{\text{id}}$. Therefore, $\mu'_{\text{id}} = \mu$ is stable and the proof is complete. \square

Corollary 6.3 (Complexity of Determining the Marital Status of a Given Participant — Lower Bound). *Theorems 3.4 and 6.1 hold also for the problem of determining whether a given participant $p \in W \cup M$ is single in some (equivalently, in every) stable marriage, where $P_W \in \mathcal{P}(W, M)$ and $P_M \in \mathcal{P}(M, W)$.*

Lemma 6.4 (Is Participant Single? \leftrightarrow Is Couple Sometimes/Always Married?). *Let $n \in \mathbb{N}$, and let W, W', M and M' be pairwise-disjoint sets, each of cardinality n ; let $w \in W$ and $m' \in M'$. There exist functions $P_{W \cup W'} : \mathcal{P}(W, M) \rightarrow \mathcal{F}(W \cup W', M \cup M')$ and $P_{M \cup M'} : \mathcal{P}(M, W) \rightarrow \mathcal{F}(M \cup M', W \cup W')$ s.t. for every $P_W \in \mathcal{P}(W, M)$ and $P_M \in \mathcal{P}(M, W)$, the following are equivalent.*

- w is single in some marriage between W and M that is stable w.r.t. P_W and P_M .
- w and m' are married in **some** marriage between $W \cup W'$ and $M \cup M'$ that is stable w.r.t. $P_{W \cup W'}(P_W)$ and $P_{M \cup M'}(P_M)$.
- w and m' are married in **every** marriage between $W \cup W'$ and $M \cup M'$ that is stable w.r.t. $P_{W \cup W'}(P_W)$ and $P_{M \cup M'}(P_M)$.

Proof. The proof is similar to that of Lemma 5.4. Denote $W = \{w_1 = w, w_2, \dots, w_n\}$, $M = \{m_1, \dots, m_n\}$, $W' = \{w'_1, \dots, w'_n\}$, and $M' = \{m'_1 = m', m'_2, \dots, m'_n\}$.

To define $P_{W \cup W'}(P_W)$, for every i we define the preference list of w_i to consist of her preference list in P_W (in the same order), followed by m'_i , followed by all other men in arbitrary order; we define the preference list of w'_i to consist of m_i , followed by all other men in arbitrary order. Similarly, to define $P_{M \cup M'}(P_M)$, for every j we define the preference list of m_j to consist of his preference list in P_M (in the same order), followed

by w'_j , followed by all other women in arbitrary order; we define the preference list of m'_j to consist of w_j , followed by all other women in arbitrary order.

Similarly to the proof of Lemma 5.4, we have that w is single in some marriage μ between W and M that is stable w.r.t. P_W and P_M iff w and m' are married in some marriage (a corresponding “supermarriage” of μ) between $W \cup W'$ and $M \cup M'$ that is stable w.r.t. $P_{W \cup W'}(P_W)$ and $P_{M \cup M'}(P_M)$. Additionally, by Theorem 2.7 (in conjunction with Theorem 2.4), we have: w is single in some marriage between W and M that is stable w.r.t. P_W and $P_M \iff w$ is single in every marriage between W and M that is stable w.r.t. P_W and $P_M \iff w$ and m' are married in every marriage between $W \cup W'$ and $M \cup M'$ that is stable w.r.t. $P_{W \cup W'}(P_W)$ and $P_{M \cup M'}(P_M)$. \square

Acknowledgements

This work was supported in part by ISF grant 230/10 and by the European Research Council under the European Community’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no. [249159].

References

- L. E. Dubins and D. Freedman. Machiavelli and the Gale-Shapley algorithm. *American Mathematical Monthly*, 88(7):485–494, 1981.
- D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, 1962.
- Y. A. Gonczarowski and E. Friedgut. Sisterhood in the Gale-Shapley matching algorithm. *The Electronic Journal of Combinatorics*, 20(2):#P12 (18pp), 2013.
- D. Gusfield. Three fast algorithms for four problems in stable marriage. *SIAM Journal on Computing*, 16(1):111–128, 1987.
- B. Kalyanasundaram and G. Schintger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5(4):545–557, 1992.
- D. E. Knuth. *Marriage stables et leurs relations avec d’autres problèmes combinatoires*. Les Presses de l’Université de Montréal, Montréal, Quebec, Canada, 1976.
- E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, Cambridge, UK, 1997.
- D. G. McVitie and L. B. Wilson. The stable marriage problem. *Communications of the ACM*, 14(7):486–490, 1971.
- C. Ng and D. S. Hirschberg. Lower bounds for the stable marriage problem and its variants. *SIAM Journal on Computing*, 19(1):71–77, 1990.
- A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, 1992.
- A. E. Roth. On the allocation of residents to rural hospitals: A general property of two-sided matching markets. *Econometrica*, 54(4):425–427, 1986.

L. B. Wilson. An analysis of the stable marriage assignment algorithm. *BIT Numerical Mathematics*, 12(4):569–575, 1972.

A. C.-C. Yao. Some complexity questions related to distributive computing (preliminary report). In *Proceedings of the 11th Annual ACM Symposium on Theory of Computing (STOC)*, pages 209–213, 1979.

A Verifying the Output of a Given Stable Marriage Mechanism

As noted in Section 3, while the lower bounds of Theorems 3.1 and 4.1 are tight, we do now know whether those of Theorems 3.3 and 5.1 are tight as well. We note that we do not even know a tight lower bound for verifying whether a given marriage is the M -optimal stable marriage.

Open Problem 2. What is the worst-case complexity of verifying whether a given marriage is the M -optimal stable marriage?

As in the case of Open Problem 1, we do not have any $o(n^2 \log n)$ algorithm for verification of the M -optimal stable marriage, even randomized and even in the strong two-party communication model, nor do we have any $\omega(n^2)$ lower bound, even for deterministic algorithms and even in the simple comparison model.

In this appendix, we derive a $\Omega(n^2)$ lower bound for verification of the M -optimal stable marriage. In fact, we show this lower bound not only for verifying the M -optimal stable marriage, but also for verifying the output of any other stable marriage mechanism.

Definition A.1 (Stable Marriage Mechanism). Let $n \in \mathbb{N}$, let W and M be disjoint sets s.t. $|W| = |M| = n$. A *stable marriage mechanism* is a function f from $\mathcal{F}(W, M) \times \mathcal{F}(M, W)$ to the set of perfect marriages between W and M , s.t. for every $P_W \in \mathcal{F}(W, M)$ and $P_M \in \mathcal{F}(M, W)$, the marriage $f(P_W, P_M)$ is stable w.r.t. P_W and P_M .

Example A.2 (M -Optimal Stable Marriage Mechanism). The function $f_{M\text{-Opt}}$, where $f_{M\text{-Opt}}(P_W, P_M)$ is the M -optimal stable marriage w.r.t. P_W and P_M , is a well-defined stable marriage mechanism by Theorem 2.4.

Corollary A.3 (Complexity of Computation of a Given Stable Marriage Mechanism — Lower Bound). *By Theorem 3.3, we have that for every stable marriage mechanism f , the worst-case query complexity (using Boolean queries as defined there) of computing f is $\Omega(n^2)$.*

Theorem A.4 (Complexity of Verification of the Output of a Given Stable Marriage Mechanism — Lower Bound). *Let $n \in \mathbb{N}$, let W and M be disjoint sets s.t. $|W| = |M| = n$, fix a stable marriage mechanism f and let $P_W \in \mathcal{F}(W, M)$ and $P_M \in \mathcal{F}(M, W)$. Let μ_{id} be the perfect marriage in which w_i is married to m_i for every i . For every type of Boolean queries that query the women’s preferences separately from the men’s preferences, every randomized (or deterministic) algorithm for determining whether $f(P_W, P_M) = \mu_{\text{id}}$ must make $\Omega(n^2)$ queries in the worst case.*

The proof of Theorem A.4 uses the machinery of Section 5, with Lemma 5.4 replaced by the following lemma.

Lemma A.5. *Let $n \in \mathbb{N}$, and let W, W', M and M' be pairwise-disjoint sets, each of cardinality n . Let μ_{id} be the perfect marriage between W and M in which w_i is married to m_i for every i , and let μ'_{id} be the perfect marriage between $W \cup W'$ and $M \cup M'$ in which for every i , both w_i is married to m_i and w'_i is married to m'_i . There exist functions $P_{W \cup W'} : \mathcal{P}(W, M) \rightarrow \mathcal{F}(W \cup W', M \cup M')$ and $P_{M \cup M'} : \mathcal{P}(M, W) \rightarrow \mathcal{F}(M \cup M', W \cup W')$ s.t. for every $P_W \in \mathcal{P}(W, M)$ and $P_M \in \mathcal{P}(M, W)$, both of the following hold.*

1. *If μ_{id} is the unique stable marriage w.r.t. P_W and P_M , then μ'_{id} is the unique stable marriage w.r.t. $P_{W \cup W'}(P_W)$ and $P_{M \cup M'}(P_M)$.*
2. *If μ_{id} is unstable w.r.t. P_W and P_M , then μ'_{id} is unstable w.r.t. $P_{W \cup W'}(P_W)$ and $P_{M \cup M'}(P_M)$.*

Proof. We define $P_{W \cup W'}(P_W)$ and $P_{M \cup M'}(P_M)$ as in Lemma 5.4, only with M' appearing sorted by j (as opposed to in arbitrary order) on the preference lists of W' , and with W' appearing sorted by i (as opposed to in arbitrary order) on the preference lists of M' . By Lemma 5.4, we have both that Part 2 holds, and that if μ_{id} is the unique stable marriage w.r.t. P_W and P_M , then it is a submarriage of every marriage that is stable w.r.t. $P_{W \cup W'}(P_W)$ and $P_{M \cup M'}(P_M)$; it is straightforward to show that every “supermarriage” of μ_{id} , apart from μ'_{id} , is unstable, thus proving Part 1 as well. \square

We prove Theorem A.4 by showing a $\Omega(n^2)$ lower bound on the randomized (and deterministic) communication complexity of verifying the output of a given stable marriage mechanism f (e.g. $f_{M\text{-Opt}}$), where P_W is held by Alice and P_M is held by Bob (and both know f); this is done similarly to the proof of Theorem 5.1. Given a protocol for verifying the output of f , a protocol for disjointness with the same communication complexity can be constructed as follows using Lemmas 5.2 and A.5: Alice computes $P_A \triangleq P_{W \cup W'}(P_W(\bar{x}))$, Bob computes $P_B \triangleq P_{M \cup M'}(P_M(\bar{y}))$, and they use the given protocol to determine whether $f(P_A, P_B) = \mu'_{\text{id}}$; by Lemmas 5.2 and A.5, we have that $\bigvee_{(i,j) \in [n]^2} x_j^i y_j^i = 0$ iff indeed $f(P_A, P_B) = \mu'_{\text{id}}$.

Open Problem 3. Is there a stable marriage mechanism whose worst-case output verification complexity is $\Theta(n^2)$? Which stable marriage mechanisms have the lowest asymptotic worst-case output verification complexity?

B Nondeterminism

All the lower bounds in this paper are based upon reductions to the well-studied communication complexity of the disjointness function. Since the disjointness function also has $\Theta(n)$ *nondeterministic* communication complexity (see Kushilevitz and Nisan, 1997), it follows that all our lower bounds apply not only to randomized communication complexity, but also to nondeterministic communication complexity. For nondeterministic communication complexity, the $\Omega(n^2)$ lower bound for finding a stable marriage is in fact tight (and so still is the $\Omega(n^2)$ bound for verification of stability).

For the decision problem of verifying the stability of a given marriage, the *co-nondeterministic* communication complexity may be easily seen to be $\Theta(\log n)$. In this appendix, we show, in contrast, a $\Omega(n^2)$ lower bound also for the *co-nondeterministic* communication complexities of determining the marital status of a given couple or participant considered in Section 6.

Theorem B.1 (Nondeterministic Communication Complexity of Determining the Marital Status of a Given Couple — Lower Bound). *In the notation of Section 3, both the nondeterministic and co-nondeterministic communication complexities of determining whether w_1 and m_1 are married in some (alternatively, in every) stable marriage, where P_W is held by Alice and P_M is held by Bob, is $\Omega(n^2)$.*

Recall that all of the proofs in Section 6 hinge on embedding disjointness in the relevant problems via the intermediate problem of whether a given participant is single in some stable marriage. In order to show the lower bounds on the co-nondeterministic communication complexities as well, it is therefore enough to embed disjointness in the complement of the problem of whether a given participant is single in some stable marriage. We do so by embedding this problem in its own complement (deducing also that a result analogous to Theorem B.1 holds for this problem as well).

Lemma B.2 (Is Participant Single? $\leftrightarrow \neg$ Is Participant Single?). *Let $n \in \mathbb{N}$, let W and M be sets s.t. $|W| = |M| = n$, and let w' and m' s.t. W , M , $\{w'\}$ and $\{m'\}$ are pairwise disjoint; let $w \in W$. There exist functions $P_{W \cup \{w'\}} : \mathcal{P}(W, M) \rightarrow \mathcal{P}(W \cup \{w'\}, M \cup \{m'\})$ and $P_{M \cup \{m'\}} : \mathcal{P}(M, W) \rightarrow \mathcal{P}(M \cup \{m'\}, W \cup \{w'\})$ s.t. for every $P_W \in \mathcal{P}(W, M)$ and $P_M \in \mathcal{P}(M, W)$, the following are equivalent.*

- w is single in some marriage between W and M that is stable w.r.t. P_W and P_M .
- m' is married in every marriage between $W \cup \{w'\}$ and $M \cup \{m'\}$ that is stable w.r.t. $P_{W \cup \{w'\}}(P_W)$ and $P_{M \cup \{m'\}}(P_M)$.

Proof. To define $P_{W \cup \{w'\}}(P_W)$, we define the preference list of w as her preference list in P_W (in the same order), followed by m' ; we define the preference list of every other woman in W as her preference list in P_W (in the same order and with m' absent), and define the preference list of w' to be empty (once again, w' can be ignored, and is defined purely for aesthetic reasons — so that $W \cup \{w'\}$ and $M \cup \{m'\}$ be of equal cardinality). To define $P_{M \cup \{m'\}}(P_M)$, we define the preference list of every man in M as his preference list in P_M (in the same order and with w' absent); we define the preference list of m' to consist solely of w .

Directly from definition of $P_{M \cup \{m'\}}$ and $P_{W \cup \{w'\}}$, we have that a natural bijection $\mu \mapsto \mu'$ exists between stable marriages w.r.t. P_W and P_M and stable marriages w.r.t. $P_{W \cup \{w'\}}(P_W)$ and $P_{M \cup \{m'\}}(P_M)$; this bijection is given by:

- If w is married in μ , then $\mu' \triangleq \mu$ (with m' and w' single in μ').
- If w is single in μ , then μ' is the marriage obtained from μ by marrying w to m' (with w' once again single in μ').

Once again by Theorem 2.7 (in conjunction with Theorem 2.4), and by the existence of this bijection, we have: w is single in some marriage between W and M that is stable w.r.t. P_W and $P_M \iff w$ is single in every marriage between W and M that is stable w.r.t. P_W and $P_M \iff m'$ is married in every marriage between $W \cup \{w'\}$ and $M \cup \{m'\}$ that is stable w.r.t. $P_{W \cup \{w'\}}(P_W)$ and $P_{M \cup \{m'\}}(P_M)$. \square

We note that the nondeterministic lower bound of $\Omega(n^2)$ for determining whether a given couple is married in some stable marriage, as well as the co-nondeterministic lower bound of $\Omega(n^2)$ for determining whether a given couple is married in every stable

marriage (and both the nondeterministic and co-nondeterministic lower bounds of $\Omega(n^2)$ for determining whether a given participant is single in some/every stable marriage), is tight. (Recall that we do not know whether any of these problems can be deterministically or even probabilistically solved using $o(n^2 \log n)$ communication.) The questions of a tight co-nondeterministic lower bound for the former problem and a tight nondeterministic lower bound for the latter remain open in all query models. We note that the latter problem may be solved by checking whether the pair in question is married in both the M -optimal stable marriage and the W -optimal stable marriage; a $O(n^2)$ -Boolean-queries algorithm (even a nondeterministic one) for verification of the M -optimal stable marriage (see Open Problem 2 in Appendix A) would therefore also settle the question of the nondeterministic communication complexity of this problem.

C Optimality of the Deferred Acceptance Algorithm w.r.t. Queries onto Women

Gale and Shapley’s (1962) proof of Theorem 2.4 is constructive, providing an efficient algorithm for finding the M -optimal stable marriage. In this algorithm, men are asked queries of the form “which woman is next on the preference list of man m after woman w ?” (or alternatively, “which woman does man m rank at place k ?”), while women are asked queries of the form “whom does woman w prefer most out of the set of men \tilde{M} ?”; all of these queries require an answer of length $O(\log n)$ bits.

Dubins and Freedman (1981) presented a variant of Gale and Shapley’s algorithm, which runs in the same worst-case time complexity, but performs a significantly-more-limited class of queries, namely only pairwise-comparison queries, onto women. In Open Problem 1 in the Introduction, we raise the question of a tight lower bound for the complexity of finding a stable marriage using only such queries *for both women and men*. In this appendix, we show that regardless of how complex the queries onto the men may be, no algorithm for finding any stable marriage (and even no algorithm for verifying the stability of a given marriage, when input a stable marriage) that performs only pairwise-comparison queries onto women, may perform any less such queries onto them than Dubins and Freedman’s variant of Gale and Shapley’s algorithm (given the same preference lists). For the duration of this appendix, let $n \in \mathbb{N}$, let W and M be disjoint sets s.t. $|W| = |M| = n$.

Definition C.1 (Pairwise-Comparison Query). A pairwise-comparison query onto W is a query of whether $m \succ_w m'$ for some given $w \in W$ and $m, m' \in M$

Definition C.2 (Men-Proposing Deferred Acceptance Algorithm (Dubins and Freedman, 1981)). The following algorithm is henceforth referred to as the *men-proposing deferred-acceptance algorithm*: The algorithm is initialized with all women and all men being *provisionally single*, and concludes when no man is provisionally single. The algorithm is divided into steps, to which we refer as *nights*. On each night, an arbitrary provisionally-single man m is chosen, and serenades under the window of the woman w ranked highest on his preference list among those who have not (yet) rejected him. If w is provisionally single, then m and w are *provisionally married* to each other. Otherwise, i.e. if w is already provisionally married to some man m' , then if $m \succ_w m'$, then w rejects m' , who becomes provisionally single, and w and m are provisionally married to each other; otherwise, w rejects m , who remains provisionally single. The algorithm stops when no

provisionally-single men remain, and the couples married by the output marriage are exactly those that are provisionally married when the algorithm stops.

Theorem C.3 (Dubins and Freedman (1981)). *Let $P_W \in \mathcal{F}(W, M)$ and $P_M \in \mathcal{F}(M, W)$ be profiles of full preference lists for W over M and for M over W , respectively. The men-proposing deferred-acceptance algorithm stops after $O(n^2)$ nights, and yields the M -optimal stable marriage.*

Remark C.4. Let $P_W \in \mathcal{F}(W, M)$ and $P_M \in \mathcal{F}(M, W)$. All runs of the men-proposing deferred acceptance algorithm (given P_W and P_M) perform the same number of pairwise-comparison queries onto W .

Theorem C.5 (Optimality of Men-Proposing Deferred-Acceptance Algorithm w.r.t. Pairwise-Comparison Queries onto W). *For any profiles $P_W \in \mathcal{F}(W, M)$ and $P_M \in \mathcal{F}(M, W)$ of full preference lists for W over M and for M over W , respectively, every algorithm for finding or verifying a stable marriage (for the latter — when input any marriage that is stable w.r.t. P_W and P_M) that only performs pairwise-comparison queries onto W (and arbitrary queries onto M), performs no less queries onto W than the men-proposing deferred-acceptance algorithm, when input P_W and P_M .*

Remark C.6. An analogous result may similarly be shown to hold w.r.t. profiles of arbitrary preference lists, and finding/verifying a possibly-imperfect stable marriage.

Definition C.7. Let μ be a perfect marriage between W and M . By slight abuse of notation, we denote the woman married to a man $m \in M$ in μ by $\mu(m)$ instead of $\mu^{-1}(m)$.

Proof of Theorem C.5. Let A be a run of the men-proposing deferred-acceptance algorithm w.r.t. P_W and P_M , and let B be a given run of an algorithm for finding/verifying a stable marriage w.r.t. P_W and P_M . Let $Q \subseteq W \times M^2$ be the set of triples (w, m, m') s.t. either the query of whether $m \succ_w m'$ was performed onto W during B and answered positively, or the query of whether $m' \succ_w m$ was performed onto W during B and answered negatively. By definition, at least $|Q|$ queries onto W are performed during B . Let μ be the M -optimal stable marriage w.r.t. P_W and P_M , i.e. the marriage output by A . Let $R \triangleq \{(w, m) \mid w \text{ rejects } m \text{ during } A\} \subseteq W \times M$. By definition, we note that the number of queries onto W during A equals the number of rejections performed during A , and so, as no woman rejects the same man twice, equals $|R|$. It is therefore enough to show that $|R| \leq |Q|$ in order to complete the proof.

Let μ' be the output of B if it is a run of an algorithm for finding a stable marriage, or the input to B if it is a run of an algorithm for verifying stability; either way, μ' a stable marriage w.r.t. P_W and P_M . We claim that $w \succ_m \mu'(m)$ for every $(w, m) \in R$. Indeed, as m serenades under women's windows during A in descending order of preference, the fact that w rejects m during A implies $w \succ_m \mu(m)$. By Theorem C.3, we thus have $w \succ_m \mu(m) \succeq_m \mu'(m)$, as claimed. As B guarantees the stability of μ' , it must therefore ascertain that $\mu'(w) \succ_w m$ for every $(w, m) \in R$; therefore, as only pairwise-comparison queries are performed onto W during B , there exists $m' \in M$ s.t. $(w, m', m) \in Q$. We have thus shown that R is contained in the projection of Q over its first and last coordinates, and therefore $|R| \leq |Q|$, and the proof is complete. \square