Stable Matchings with Diversity Constraints: Affirmative Action is beyond NP

Jiehua Chen, **Robert Ganian** and **Thekla Hamm**
TU Wien, Vienna, Austria
jiehua.chen@tuwien.ac.at, rganian@gmail.com, thamm@ac.tuwien.ac.at

Abstract

We investigate the following many-to-one stable matching problem with diversity constraints (SMTI-DIVERSE): Given a set of students and a set of colleges which have preferences over each other, where the students have overlapping types, and the colleges each have a total capacity as well as quotas for individual types (the diversity constraints), is there a matching satisfying all diversity constraints such that no unmatched student-college pair has an incentive to deviate?

SMTI-DIVERSE is known to be NP-hard. However, as opposed to the NP-membership claims in the literature [Aziz *et al.*, 2019; Huang, 2010], we prove that it is beyond NP: it is complete for the complexity class Σ_2^{P}. In addition, we provide a comprehensive analysis of the problem's complexity from the viewpoint of natural restrictions to inputs and obtain new algorithms for the problem.

1 Introduction

Stability is a classic and central property of assignments, or *matchings*, of agents to each other, describing that no two agents prefer each other to their respective situations in the matching. Stability is desirable in many scenarios and spawned numerous works in various contexts [Manlove, 2013]. In this work we investigate the notion of stability in combination with diversity, which is key in many real-world matching applications, ranging from education, through health-care systems, to job and housing markets [Abdulkadiroğlu, 2005; Huang, 2010; Kamada and Kojima, 2015; Kurata *et al.*, 2017; Ahmed *et al.*, 2017; Benabbou *et al.*, 2019; Gonczarowski *et al.*, 2019; Aziz *et al.*, 2019].

For this we conceptually distinguish two sets—a set of *students* which should be matched to a set of *colleges* (each with a maximum *capacity* to accommodate students) with the additional constraint that the set of students matched to any single college has to be diverse. The diversity requirements are modeled as *types* which are attributes that a student may or may not have, and *upper and lower quotas* that specify how many students of a certain type may be matched to a given college. The terminology arises from the context of controlled public school choice, a typical application of this paradigm where it is desirable to match colleges to students to ensure stability as well as demographic, socio-economic, and ethnic diversity (see also *affirmative action*).

The study of stable matchings with diversity constraints was initiated by Abdulkadiroğlu [2005] in the context of college admissions. It has since become an ongoing and actively researched topic among economists and computer scientists, covered for example by two chapters [Heo, 2019; Kojima, 2019] in the recently published book "On the Future of Economic Design" [Laslier *et al.*, 2019]. One of the fundamental questions in this area is whether there is a diverse and stable matching between students and colleges; the corresponding computational problem is called SMTI-DIVERSE (see Section 2 for formal definitions).

As has already been observed in the work of Aziz *et al.* [2019] and hinted at in Huang's earlier work on a closely related problem [2010], SMTI-DIVERSE is NP-hard. The authors further claimed that the problem(s) under consideration belong to NP (see also [Manlove, 2013, Chapter 5.2.5]). We disprove this claim by presenting an involved reduction showing that the problem is complete for the complexity class Σ_2^{P}, even under severe restrictions to the input instances. Hence, the problem is substantially more difficult than all NP problems, unless a widely believed complexity-theoretical assumption collapses. In particular, this implies that the problem is not easily amenable to SAT or ILP solvers.

Complementing this hardness finding, we systematically analyze the complexity of the problem by considering natural relaxations (such as dropping lower quotas or dropping stability) or restrictions (such as bounding the number n of students, the number t of types, the number m of colleges, and/or the maximum upper quota u_∞, and the maximum capacity q_∞). The outcome of our analysis is a full classification of the complexity of SMTI-DIVERSE with respect to the considered restrictions and relaxations, presented in Table 1. We highlight three key technical contributions of our work:

(1) SMTI-DIVERSE is Σ_2^{P}-complete even when the preferences do not have ties and $m = 4$, while two natural relaxations of the problem (either dropping the lower quotas or the stability requirement) lower the complexity to NP-complete.
(2) When n, $m+t$, or $m+\mathsf{q}_\infty$ is a constant, SMTI-DIVERSE can be solved in polynomial time.
(3) SMTI-DIVERSE is NP-complete even if lower quotas are

all zero and $t + \mathsf{u}_\infty + \mathsf{q}_\infty$ is a constant. This result also fixes a technical flaw in [Aziz *et al.*, 2019, Proposition 5.3].

Related work. For one type, SMTI-DIVERSE is equivalent to the Hospitals/Residents with Lower Quotas problem where no hospital is allowed to be closed (HR-LQ-2), as studied by Hamada *et al.* [2016]. This problem is polynomial-time solvable when no ties are allowed [Manlove, 2013, Chapter 5.2.3]. We show that SMTI-DIVERSE becomes NP-hard even for only two types.

Huang [2010] introduced the closely related CLASSIFIED STABLE MATCHING (CSM) problem, which asks for a matching that fulfills the diversity constraints and does not admit *blocking coalitions*. We show that our Σ_2^{P}-hardness reduction can be adapted to show Σ_2^{P}-completeness for CSM.

Aziz *et al.* [2019] studied school choice with diversity constraints, but with a slightly different stability condition: an unmatched student-college pair $\{u, w\}$ is *d-blocking* if it is a blocking pair (in our sense) and the new solution fulfills the diversity constraints for *all* colleges instead of only for w. This means that a d-blocking pair is also a blocking pair, but the converse is not true. However, dropping the lower quotas requirements renders both concepts equivalent. Our model of blocking pairs is a direct extension of HR-LQ-2, where a student and a college already form a blocking pair once the new solution is better for them, regardless of the other colleges' lower quotas. Such a model assumes that the blocking condition is tested based on local information of whether the deviating college's diversity constraints are fulfilled after the rematching. This is a standard assumption in many controlled school choice articles [Abdulkadiroğlu, 2005; Kurata *et al.*, 2017; Hamada *et al.*, 2016]. Nevertheless, our Σ_2^{P}-hardness reduction establishes the same hardness for their variant. Other related work includes recent papers by Nguyen and Vohra [2019], Kurata *et al.* [2017], Ismaili *et al.* [2019].

Problems	FI-DIVERSE	SMTI-DIVERSE	
Constraints	$(\ell_\infty \geq 0)$	$(\ell_\infty \geq 0)$	$(\ell_\infty = 0)$
Complexity	NP-c$^\diamond$	Σ_2^{P}-c [Th 1]	NP-c [Th 2]
$m + \mathsf{u}_\infty$	NP-c$^\diamond$	Σ_2^{P}-c [Th 1]	NP-c [Th 3]
$t + \mathsf{u}_\infty + \mathsf{q}_\infty$	NP-c [Pr 2]	NP-c$^\clubsuit$ [Th 2,Ob 1]	NP-c [Th 2]
n	P [Th 4]	P [Th 4]	P [Th 4]
$m + t$	P [Th 5]	P [Th 5]	P [Th 5]
$m + \mathsf{q}_\infty$	P [Pr 3]	P [Pr 3]	P [Pr 3]

Table 1: A complete picture of the complexity results for FI-DIVERSE and SMTI-DIVERSE (see Section 2 for the definitions). Results marked with $^\diamond$ are due to [Aziz *et al.*, 2019, Prop 5.1] while the remaining ones are new. All hardness results hold even for preferences with *no* ties, even if the corresponding measures are upper-bounded by a constant. The NP-containment result marked with $^\clubsuit$ holds already when either t or q_∞ is a constant (see Observation 1).

2 Preliminaries

Given an integer z, let $[z]$ denote the set $\{1, \ldots, z\}$. Given two integer vectors $\boldsymbol{x}, \boldsymbol{y}$ of the same dimension, i.e., $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Z}^z$ for a non-negative integer z, we write $\boldsymbol{x} \leq \boldsymbol{y}$ if for each index $i \in [z]$ it holds that $\boldsymbol{x}[i] \leq \boldsymbol{y}[i]$; otherwise, we write $\boldsymbol{x} \not\leq \boldsymbol{y}$. A *preference list* $\succeq$ over a set A is a complete and transitive binary relation on A. We use $\succ$ to denote the asymmetric part (i.e., $x \succeq y$ and $\neg(y \succeq x)$) and $\sim$ the symmetric part of $\succeq$ (i.e., $x \succeq y$ and $y \succeq x$). We say that x is *(strictly) preferred* (resp. *weakly preferred*) to y if $x \succ y$ (resp. $x \succeq y$), and that x and y are *tied* in $\succeq$ if $x \sim y$; $\succeq$ is said to *contain ties* in this case. We write $[A]$ to denote an arbitrary but fixed linear order on A. The expression "$x \succ Y$" (resp. "$x \succeq Y$") means that x is strictly (resp. weakly) preferred to every one in Y.

Problem-specific terminology. The problem we study has as input a set $T := [t]$ of types, a set $U := \{u_1, u_2, \ldots, u_n\}$ of n students and a set $W := \{w_1, w_2, \ldots, w_m\}$ of m colleges together with the following information. Each student $u \in U$ has (i) a *preference list* $\succeq_u$ over a subset $\mathsf{A}(u) \subseteq W$ of the colleges, and (ii) a *type vector* $\tau_u \in \{0, 1\}^t$, where $\tau_u[z] = 1$ means that u has type z. Each college $w \in W$ has (i) a *preference list* $\succeq_w$ over a subset $\mathsf{A}(w) \subseteq U$ of the students, and (ii) a *lower-quota* and *upper-quota* for each type which is described, respectively, via the vectors ℓ_w and $\mathsf{u}_w \in [n]^t$, where $\ell_w \leq \mathsf{u}_w$, and (iii) a *capacity* $q_w \in [n]$ which is the maximum number of students allowed to be admitted to w. Note that while the capacity can be modeled by introducing an extra type, separating it from the types allows for a more refined analysis of the problem's complexity.

For each $x \in U \cup W$, we call $\mathsf{A}(x)$ the *acceptable set* of x, which contains all students or colleges that are acceptable to x. Throughout the paper, we assume that no student or college has an empty acceptable set, and for each student u and each college w it holds that $u \in \mathsf{A}(w)$ iff. $w \in \mathsf{A}(u)$.

A *matching* M is a set of student-college pairs of the form $\{u, w\}$, where each student u is involved in *at most one* pair in M and $w \in \mathsf{A}(u)$. If $\{u, w\} \in M$, then we say that u and w are *assigned* to each other by M. Slightly abusing the notation, given a student $u \in U$ if there exists a college $w \in W$ with $\{u, w\} \in M$, then define $M(u) := w$; otherwise define $M(u) := \bot$. We assume that each student u prefers every acceptable college $w \in \mathsf{A}(u)$ to $\bot$. Similarly, given a college $w \in W$, we write $M(w) := \{u \mid \{u, w\} \in M\}$ to denote the set consisting of all students assigned to w by M.

Feasible and stable matchings. A matching M is *feasible* for an instance $(U, W, T, (\tau_u, \succeq_u)_{u \in U}, (\succeq_w, \ell_w, \mathsf{u}_w, q_w)_{w \in W})$ if it is *feasible* for each college $w \in W$, i.e., college w (i) is assigned at most q_w students, i.e., $|M(w)| \leq q_w$, and (ii) meets the lower and upper quotas for each type, i.e., $\ell_w \leq \sum_{u \in M(w)} \tau_u \leq \mathsf{u}_w$.

A student u and a college w form a *blocking pair* in a matching M if: (i) $u \in \mathsf{A}(w)$ and $\{u, w\} \notin M$, (ii) $w \succ_u M(u)$, (iii) there exists a (possibly empty) subset $U' \subseteq M(w)$ such that $u \succ_w U'$, and (iv) $M \cup \{\{u, w\}\} \setminus \{\{u', w\} \mid u' \in U'\}$ is feasible for w. Accordingly, we say that U' is a *witness* for $\{u, w\}$ to block M. A matching M is *stable* if it has *no* blocking pairs.

Problem variants. Now, we formally state our main problem of interest—the natural generalization of the classical MANY-TO-ONE STABLE MATCHING WITH TIES AND IN-

COMPLETE PREFERENCES (SMTI) [Manlove, 2013] to incorporate diversity constraints:

SMTI-DIVERSE
Input: A set U of n students, a set W of m colleges, a set T of types, the type vectors and preference lists $(\tau_u, \succeq_u)_{u\in U}$ for the students, the preference lists, lower-quota vectors, upper-quota vectors, and capacities $(\succeq_w, \ell_w, \mathsf{u}_w, q_w)_{w\in W}$ for the colleges.
Question: Is there a feasible and stable matching?

We use SMI-DIVERSE to denote the restriction of SMTI-DIVERSE to the case with *no* ties. Moreover, we use FI-DIVERSE to denote the problem of deciding whether there is a feasible matching (representing a generalization for FEASIBLE MATCHING WITH INCOMPLETE PREFERENCES).

For an illustration, let there be four students $u_1, \ldots, u_4$, two colleges w_1, w_2, with the following type vectors (T.) and the preference lists (Pref.) of the students (S.) as well as the preference lists (Pref.), the lower quotas (LQ.), the upper quotas (UQ.), and the capacities (C.) of the colleges depicted as follows; **preferences always ordered by** $\succ$:

S.	Pref.	T.	S.	Pref.	T.	C.	Pref.	LQ.	UQ.	C.
u_1:	$w_1\, w_2$	01	u_2:	$w_1\, w_2$	01	w_1:	$u_3\, u_2\, u_1$	11	22	2
u_3:	$w_2\, w_1$	11	u_4:	w_2	10	w_2:	$u_1\, u_3\, u_4\, u_2$	11	11	2

There are two feasible matchings M_1 and M_2 with $M_1(w_1)=\{u_1, u_3\}$, $M_1(w_2)=\{u_2, u_4\}$, $M_2(w_1)=\{u_2, u_3\}$, $M_2(w_2)=\{u_1, u_4\}$. But M_1 is blocked by $\{u_3, w_2\}$ while M_2 is stable. If u_1 does not accept w_2, then no feasible and stable matching exists.

3 How Hard is Diversity?

General complexity. Aziz *et al.* [2019, Proposition 5.1] proved that FI-DIVERSE is NP-complete; the hardness result holds even for a single college. They also claimed that determining a matching without any d-blocking pairs is NP-complete [2019, Proposition 5.3]. However, the proof used to show NP-membership is technically flawed—in particular, while the proof claims that "*Deciding whether a stable outcome exists is in NP, since we can guess an outcome X and check whether X admits blocking pair in polynomial time*", by adapting the reduction in [Aziz *et al.*, 2019, Proposition 5.1] we can show that this is impossible unless coNP ⊆ P.

Proposition 1. *Deciding whether a given feasible matching has no blocking pairs or no d-blocking pairs is* coNP*-hard.*

Note that Proposition 1 itself does not rule out that SMI-DIVERSE is in NP. It just suggests the given proof is incorrect. There could in principle be a different non-deterministic algorithm to place the problem in NP. We show that this is not the case in our main result (Theorem 1), by showing Σ_2^P-hardness. For this we introduce a crucial gadget which is used in several of our reductions throughout this section.

Lemma 1. *Let $T = \{1, 2\}$, and let $U \uplus \{r_1, r_2, r_3\}$ be a set of students with three distinguished students r_1, r_2, r_3, and let $W \uplus \{a, b\}$ be a set of colleges with two distinguished colleges a and b. Similar to the format given in Section 2, the preference lists and type vectors of students r_1, r_2, and r_3, and the preference lists, the upper quotas, and the capacities of the colleges are depicted as follows:*

S.	*Pref.*	*T.*	*C.*	*Pref.*	*UQ.*	*C.*
r_1:	$b\, a$	10	$\forall w \in W$:	$[U]\, r_2$	11	q_w
r_2:	$b\, [W]\, a$	11	a:	$r_1\, r_2\, r_3$	11	1
r_3:	$a\, b$	01	b:	$r_3\, r_2\, r_1$	11	2

All students in U have zero types and arbitrary but fixed preferences. All lower quotas are zero. The following holds for every matching M. (1) If $M(a)=\{r_2\}$, $M(b)=\{r_1, r_3\}$, and $|M(w) \cap U|=q_w$ for all $w \in W$, then no *pair $\{u, w\}$ with "$u\in\{r_1, r_2, r_3\}$ and $w\in\{a, b\}$" or with "$u = r_2$ and $w \in W$" is blocking M. (2) If $|M(w) \cap U|<q_w$ for some $w\in W$, then M is not stable.*

It is straightforward to verify the correctness of Lemma 1 by case analysis. With Lemma 1 in hand, we can prove that the problem is not NP-complete but instead lies on the second level of the polynomial hierarchy.

Theorem 1. SMTI-DIVERSE *is Σ_2^{P}-complete, and remains Σ_2^{P}-hard even if feasible matchings always exist, there are no ties, $m = 4$, and $\mathsf{u}_\infty = 3$.*

Proof Sketch. To show that SMTI-DIVERSE is in Σ_2^{P}, observe that checking whether a matching is *not* stable can be done by an NP-oracle (guess an unmatched pair $\{u, w\}$ and a subset S of students, and check in polynomial time whether S witnesses that $\{u, w\}$ is blocking M). Hence, we can guess in polynomial time a matching, ask the NP-oracle whether it is not stable, and return yes iff. the oracle answers no. Containment follows because $\mathrm{NP}^{\mathrm{NP}} = \Sigma_2^{\mathrm{P}}$ [Papadimitriou, 1994].

To show Σ_2^{P}-hardness, we reduce from a problem called NOT-1-IN-3-∃∀3SAT: Given a Boolean 3CNF formula $\phi(X, Y)$ over two equal-size variable sets X, Y such that each clause contains at least two literals from $Y \cup \overline{Y}$, is there a truth assignment of X such that for each truth assignment of Y at least one clause C_j is *not 1-in-3-satisfied*, i.e., C_j does not have precisely one true literal. NOT-1-IN-3-∃∀3SAT can be shown to be Σ_2^{P}-hard via a standard (and complementing) reduction from ∀∃3SAT, a classic Π_2^{P}-hard problem [Stockmeyer, 1976].

The idea of our main reduction is to construct, from an instance I of NOT-1-IN-3-∃∀3SAT with $|X| = |Y| = r$ and s clauses, an equivalent instance I' of SMI-DIVERSE with $2r$ "variable-types", s "clause-types" and 2 auxiliary types (the types are ordered in this sequence). Instance I' contains a special student d that has all variable-types and all clause-types, and two distinguished colleges v, w which can both accommodate d, but d prefers being in v. I' furthermore uses Lemma 1 to construct a gadget which ensures that a matching can only be stable if d is matched to w—in particular, this will force a stable and feasible matching to ensure $\{d, v\}$ will not form a blocking pair.

Moreover, I' contains one clause-student d_j for each clause C_j (let D denote the set of all clause-students) and one student lit for each literal in $X \cup Y \cup \overline{X} \cup \overline{Y}$. Student d_j only has one type: the clause-type $2r + j$ corresponding to C_j. Student lit has the variable-type $i \in [r]$ corresponding to its variable as well as all the clause-types $2r + j$ of every clause C_j containing lit. All Y-literal students only want to go to v; all positive X-literal students prefer v to b while all negative X-literal students prefer b to v.

We can now explain the core of the reduction: the quotas of v are set up in a way which ensures (assuming d is matched to w) that precisely one literal-student for each variable in X, both literal-students for each variable in Y, and some clause-students must be matched to v. In particular, a clause-student d_j will be matched to v if and only if the literal-student missing from v represents a literal in C_j. Once set up, we show that $\{d, v\}$ is blocking if and only if there is a witness set of literal-students, and this witness set would represent an assignment which 1-in-3 satisfies I. In other words, a feasible and stable matching M exists if and only if there is an assignment of the X-variables (which can be reconstructed from M) such that no assignment of the Y-variables 1-in-3-satisfies all clauses.

The following describes the preference lists, quotas and capacities of the colleges, together with the preference lists of the students from $\{r_1, r_2, r_3, d\}$.

S. Pref.	C. Pref.	LQ.	UQ.	C.
r_1:$b\,a$	w: $d\,r_2$	0^{2r+s+2}	1^{2r+s+2}	1
r_2:$b\,w\,a$	a: $r_1\,r_2\,r_3$	0^{2r+s+2}	$0^{2r+s}11$	1
r_3:$a\,b$	b: $r_3\,r_2\,r_1\,[X]\,[\overline{X}]$	1^r0^{r+s+2}	$1^r0^s1^{s+2}$	$r+2$
d: vw	v: $[D]\,d\,[Y]\,[\overline{Y}]\,[\overline{X}]\,[X]$	$1^r2^r3^s00$	$1^r2^r3^s00$	$3r+s$

This completes the construction which can be verified to fulfill the restriction stated in the theorem. It now suffices to show that "there exists an X-assignment σ_X such that for each Y-assignment σ_Y at least one clause is *not* 1-in-3-satisfied" if and only if "matching $X' \cup Y \cup \overline{Y} \cup D'$ to v, $(\{r_1, r_3\} \cup X \cup \overline{X}) \setminus X'$ to b, r_2 to a, and d to w is feasible and stable, where X' corresponds to the assignment σ_X and $D' = \{d_j \mid |(X' \cup Y \cup \overline{Y}) \cap C_j| = 2\}$". ☐

To see why the reduction behind Theorem 1 can be used to directly show Σ_2^{P}-hardness for the problem studied by Aziz *et al.* [2019] we observe that in the constructed instance, $\{d, v\}$ is a blocking pair if and only if it is a d-blocking pair. The proof of Theorem 1 can also be adapted to correct an erroneous theorem pertaining to a related problem called CLASSIFIED STABLE MATCHING (CSM) [Huang, 2010, Theorem 3.1]. In particular, that theorem claims that CSM is NP-complete, but it is in effect also Σ_2^{P}-hard. The idea for the adaption is to construct dummy variable-students with zero-types and introduce additional types to ensure that $\{v, d\}$ is a blocking pair in the proof of Theorem 1 if and only if v forms with d and the dummy variable-students a blocking coalition.

The impact of diversity. The fact that SMI-DIVERSE lies in a higher complexity class than FI-DIVERSE can be attributed to the stability constraints. On the other end of the spectrum, a stable matching without diversity constraints always exists and can be found in polynomial time [Manlove, 2013, Chapter 3]. We can pinpoint the cause of this jump in complexity more precisely to the existence of lower quotas, which in some sense implement affirmative action in the SMI-DIVERSE model. Specifically, we show that if the lower quotas are all zero, then SMTI-DIVERSE becomes NP-complete (Theorem 2). Lemma 2 will be crucial for showing NP-containment.

Lemma 2. *If $\ell_\infty = 0$, then checking whether a matching is stable can be done in $\mathcal{O}(n \cdot m \cdot t)$ time.*

Proof Sketch. It suffices to show that a matching M is stable for an instance $I = (U, W, T, (\tau_u, \succeq_u)_{u \in U}, (\succeq_w, \ell_w = \mathbf{0}, \mathsf{u}_w, q_w)_{w \in W})$ if and only if the following (polynomially verifiable) condition is met: for each unmatched student-college pair $\{u, w\} \notin M$ with u preferring w to $M(u)$ either "$\tau_u + \sum_{u' \in M(w)\colon u' \succeq_w u} \tau_{u'} \not\leq \mathsf{u}_w$" or "$|M(w)|=q_w$ and $M(w) \succeq_w u$", or both holds. ☐

Even though with zero lower quotas SMTI-DIVERSE is in NP, and thus can be considered significantly easier than SMTI-DIVERSE in general, it is actually hard within NP. Hardness for this case was claimed in [Aziz *et al.*, 2019, Proposition 5.3]. However the reduction contains a technical flaw. Indeed, the instance constructed in that proof is always a yes instance, independent of the original 3-SAT instance. To see this, define the matching M for their produced instance as follows (notations taken from that proof): First, let $M_{\mathrm{F}} := \bigcup_{i \in [k]} X_F^i \setminus \{(t_1^i, o(t_1^i)), (t_2^i, o(t_2^i)) \mid i \in [k]\}$. For each $j \in [l]$, let S_j be the set consisting of the first two (if there are fewer than two, then all) students of the form t_z^i ($i \in [k]$, $z \in [2]$) appearing in the preferences of o_j. Then $M := M_{\mathrm{F}} \cup \{(s, o_j) \mid j{\in}[l] \wedge s{\in}S_j\}$ is feasible and stable.

Below, we use Lemma 1 to provide a new and simpler NP-hardness proof for the case with zero lower quotas.

Theorem 2. *For $\ell_\infty = 0$,* SMI-DIVERSE *is* NP-*complete; it remains* NP-*hard even if $\ell_\infty = 0$, $\mathsf{u}_\infty = 1$, and $t = \mathsf{q}_\infty = 2$.*

Proof Sketch. To show NP-containment, we guess in polynomial time a matching M, and check whether M is feasible and stable in polynomial time, using Lemma 2.

To establish NP-hardness, we reduce from (2,2)-3SAT, an NP-complete variant [Berman *et al.*, 2003] of 3SAT where each literal $\mathsf{lit} \in X \cup \overline{X}$ appears precisely two times in the set $\phi(X)$ of clauses. Given an instance $I = (X = \{x_1, \ldots, x_r\}, \phi(X) = \{C_1, \ldots, C_s\})$ of (2,2)-3SAT, construct an instance of SMI-DIVERSE as follows. For each clause $C_j \in \phi(X)$, introduce a *clause-college* c_j. For each variable $x_i \in X$, introduce two *variable-students* x_i and y_i, four *literal-students* u_i^1, u_i^2, v_i^1, and v_i^2, and two *variable-colleges* w_i and p_i. Introduce three special students r_1, r_2, r_3, and two special colleges a and b. Let $T = \{1, 2\}$.

For ease of description we use the following notation: let $c[u_i^z]$ and $c[v_i^z]$, ($z \in [2]$) be the clause-college c_j such that clause C_j contains the z^{th} occurrence of literal x_i, and $\overline{x}_i$ respectively. Further, let $s^z[c_j]$ *(*$z \in [3]$*)* denote the literal-student that corresponds to the z^{th} literal appearing in clause C_j. For instance, if $C_j = (\overline{x}_2, x_3, \overline{x}_5)$ and the occurrence of $\overline{x}_2$ in c_j is its second one, then $s^1[c_j] = v_2^2$. Types and preference lists are given below, in a format similar to the one in Lemma 1.

All lower quotas are zero. This completes the construction of the instance for SMI-DIVERSE. One can verify the restrictions stated in the theorem. We show that $(X, \phi(X))$

is satisfiable if and only if the constructed instance admits a feasible and stable matching.

S.	Pref.	T.	S.	Pref.	T.	C.	Pref.	UQ.	C.
u_i^1:	$w_i\, c[u_i^1]$	11	r_1:	$b\, a$	10	c_j:	$s^1[c_j]\, s^2[c_j]\, s^3[c_j]\, r_2$	11	1
u_i^2:	$w_i\, c[u_i^2]$	00	r_2:	$b\, [C]\, a$	11	a:	$r_1\, r_2\, r_3$	11	1
v_i^1:	$p_i\, c[v_i^1]$	11	r_3:	$a\, b$	01	b:	$r_3\, r_2\, r_1$	11	2
v_i^2:	$p_i\, c[v_i^2]$	00	x_i:	$p_i\, w_i$	10	w_i:	$x_i\, u_i^1\, y_i\, u_i^2$	11	2
			y_i:	$w_i\, p_i$	01	p_i:	$y_i\, v_i^1\, x_i\, v_i^2$	11	2

The "only if" part: let the truth assignment σ_X satisfy $\phi(X)$. It can be verified that the following matching M is feasible. • For each $x_i \in X$, if $\sigma_X(x_i) = \texttt{true}$, then let $M(w_i) := \{x_i, y_i\}$ and $M(p_i) := \{v_i^1, v_i^2\}$; otherwise let $M(w_i) := \{u_i^1, u_i^2\}$ and $M(p_i) := \{x_i, y_i\}$. • For each clause $C_j \in \phi(X)$, let $M(c_j) := \{s^z(c_j)\}$, where $z \in [3]$ is minimal such that the z^{th} literal in C_j is set to $\texttt{true}$ under σ_X; note that there exists at least one such literal since σ_X is a satisfying assignment. • Let $M(a) := \{r_2\}$ and $M(b) := \{r_1, r_3\}$.

To show that M is stable, consider a blocking student-college pair $\{\alpha, \beta\}$ for M witnessed by $S' \subseteq M(\beta)$. By Lemma 1(1), we infer that α must lie in $\{x_i, y_i, u_i^1, u_i^2, v_i^1, v_i^2\}$ for some $i \in [r]$. It then suffices to do a case distinction that rules out α being one of the former 2 students, and also being one of the latter 4 students.

For the "if" part, let M be a feasible and stable matching for the constructed SMI-DIVERSE instance. Define the following assignment σ_X with $\sigma_X(x_i) := \texttt{true}$ if there exists a clause-college c_j such that u_i^1 or u_i^2 is assigned to c_j; let $\sigma_X(x_i) := \texttt{false}$ if there exists a clause-college c_j such that v_i^1 or v_i^z is assigned to c_j. If no such clause-college exists, then the truth value of x_i can be arbitrary; e.g., let $\sigma_X(x_i) = \texttt{true}$. The constructed assignment satisfies all clauses because of Lemma 1(2). Thus, to complete the proof, it suffices to show that σ_X is a valid truth assignment. □

We note that the reduction behind Theorem 2 can be adapted to show NP-hardness for FI-DIVERSE, even with three types.

Proposition 2. FI-DIVERSE *remains* NP-*hard even if* $t = 3$, $\mathsf{u}_\infty = 1$, $\mathsf{q}_\infty = 2$.

As a final remark on the impact of diversity, we note that if there are only few types t or the maximum capacity q_∞ is a constant, then SMTI-DIVERSE is in NP. The reason for this is that the size of a witness set is upper-bounded by $\min\{t, \mathsf{q}_\infty\}$.

Observation 1. *If t or q_∞ is a constant, then* SMTI-DIVERSE *is in* NP.

The case with few colleges. The NP-hardness reduction behind Theorem 2 produces a college gadget for each variable in order to maintain as few types as possible. This leads to the question of whether the problem remains NP-hard for few colleges. The following theorem answers the question affirmatively. The idea is to reduce from the NP-complete INDEPENDENT SET problem [Garey and Johnson, 1979] and introduce types corresponding to the vertices and the edges in an input graph, and students corresponding to the vertices such that the students assigned to a special college w must correspond to an independent set. We use Lemma 1 to enforce that w receives at least some given number of students.

Theorem 3. SMI-DIVERSE *is* NP-*hard even if* $m = 4$, $\ell_\infty = 0$ *and* $\mathsf{u}_\infty = 2$.

4 Algorithmic Results

This section provides the algorithmic results that together allow us to complete Table 1. The first result deals with the case where the number of students is bounded by a constant.

Theorem 4. SMTI-DIVERSE *can be solved in* $\mathcal{O}(n \cdot m \cdot t + 2^n \cdot (2n+1)^n \cdot n^2 \cdot t)$ *time.*

Proof Sketch. We show how to preprocess an SMTI-DIVERSE instance $I = (U, W, T, (\succeq_u, \tau_u)_{u \in U}, (q_w, \ell_w, \mathsf{u}_w)_{w \in W})$ to obtain an instance $I' = (U, W', T', (\succeq'_u, \tau'_u)_{u \in U}, (q'_w, \ell'_w, \mathsf{u}'_w)_{w \in W})$ with n students, $n^2 + n$ colleges and 2^n types which is equivalent in terms of the existence of a feasible and stable matching. It then suffices to solve I' in the claimed running time via an exhaustive brute-force procedure.

2^n types. Observe that types $z, z' \in T$ which describe the same subset of students, i.e., $\{u \in U \mid \tau_u[z]=1\} = \{u \in U \mid \tau_u[z']=1\}$, can be merged into a single type ζ. For each feasible matching of I, the students assigned to a college $w \in W$ of two types $z, z' \in T$ merged in this way always adhere to the stricter of the upper and lower quotas of the merged types, i.e., $\max\{\ell_w[z], \ell_w[z']\}$ and $\min\{\mathsf{u}_w[z], \mathsf{u}_w[z']\}$. Exhaustive merging yields the modified types T' with $|T'| \leq 2^n$.

$n^2 + n$ colleges. First, note that we can reject an instance with more than n colleges with non-zero lower quotas. To upper-bound the number of colleges with zero lower-quotas, denoted as W_0, note that in a stable and feasible matching M every student (say, u) matched to a college from W_0 may only be matched to one of her n most preferred colleges in W_0 which has enough upper-quotas to accommodate her. Otherwise there would be an empty zero lower-quota college in W_0 which u prefers to $M(u)$, forming a blocking pair.

We employ this observation by defining the following marking procedure. Let us begin by setting $W' = \emptyset$. Now, for each student $u \in U$, we mark the n most preferred colleges in $W_0 \setminus W'$, resolving ties arbitrarily. Clearly, at the end we obtain a set W' of size at most n^2. This is easy to prove via a replacement argument and the above observation that the colleges in $W_0 \setminus W'$ may be deleted without changing the existence of a stable and feasible matching. □

Next, we show that SMTI-DIVERSE can be solved in polynomial time if the number m of colleges and the maximum capacity q_∞ of all colleges are constants, using a simple brute-forcing algorithm based on the following observation.

Observation 2. *Every feasible matching can assign colleges to at most $m \cdot \mathsf{q}_\infty$ students.*

By the above observation, we only need to guess a subset of at most $m \cdot \mathsf{q}_\infty$ students which are assigned to colleges, and branch for each student in the guessed set on the choice of one out of m possible colleges. For each branch, we check feasibility and stability in $\mathcal{O}(2^{\mathsf{q}_\infty} \cdot n \cdot m \cdot t)$ time since each college obtains at most q_∞ students (see Observation 1).

Proposition 3. SMTI-DIVERSE *and* FI-DIVERSE *can be solved in* $\mathcal{O}(n^{m \cdot \mathsf{q}_\infty} \cdot (m \cdot \mathsf{q}_\infty)^m \cdot 2^{\mathsf{q}_\infty} \cdot n \cdot m \cdot t)$ *time.*

Finally, we turn our attention to instances with a small number of colleges and types, and show that in this case SMTI-DIVERSE also admits a polynomial-time algorithm.

We note that while under such restrictions one can use the *bounded-variable ILP Encoding* technique [Bredereck *et al.*, 2014, Section 3.1] to show that FI-DIVERSE becomes polynomial-time solvable, the same technique is unlikely to work for SMTI-DIVERSE. That is because two students, even with the same type vectors and the same preferences, may be preferred differently by a college.

However, to witness a blocking pair in a matching, only those students assigned to a college w need to be considered that are least preferred by w among all students assigned to w with the same type vector. We introduce two notations to formally describe such students. Given a matching M, a college v, and a type vector $\tau \in \{0,1\}^t$, let $\mathsf{S}(M,v,\tau) \coloneqq \{u \in M(v) \mid \tau_u = \tau\}$ denote the set of students with type vector τ that are assigned to v, and let $\mathsf{worst}(M,v,\tau)$ denote the set of students in $M(v)$ with type vector τ that v prefers least:

$$\mathsf{worst}(M,v,\tau) \coloneqq \{u \in \mathsf{S}(M,v,\tau) \mid \mathsf{S}(M,v,\tau) \succeq_v u\}.$$

Proposition 4. *Let M be a feasible matching in an* SMTI-DIVERSE *instance. Then, an unmatched student-college pair $\{u,w\}$ with $w \succ_u M(u)$ is blocking M if and only if there is a subset of k students $U' \coloneqq \{u_{i_1},\ldots,u_{i_k}\} \subseteq M(w)$ ($0 \le k \le |M(w)|$) assigned to w such that*

(i) no two students from U' have the same type vector,
(ii) each student $u' \in U'$ belongs to $\mathsf{worst}(M,w,\tau_{u'})$,
(iii) w strictly prefers u to each student in U', and
(iv) $M \cup \{\{u,w\}\} \setminus (\{\{u,M(u)\}\} \cup \{\{u',w\} \mid u' \in U'\})$ is feasible for w.

Proof Sketch. (iii) and (iv) necessarily hold for every set that witnesses that $\{u,w\}$ is a blocking pair. (i) can be seen to hold for every minimal such set because at most one student from each type vector has to be removed from $M(w)$ to make the addition of u to $M(w)$ feasible for w. We can achieve (ii) by swapping each student $u \notin \mathsf{worst}(M,w,\tau_u)$ from a minimal witnessing set for a student in $\mathsf{worst}(M,w,\tau_u)$. This modification maintains all previous conditions and the fact that the set witnesses that $\{u,w\}$ is a blocking pair. ☐

Theorem 5. SMTI-DIVERSE *can be solved in $\mathcal{O}(n^{m\cdot 2^t+(2m+1)\cdot(t+1)} \cdot m^2 \cdot (n^t \cdot t + m))$ time.*

Proof Sketch. Let $I = (U,W,T=[t],(\tau_u,\succeq_u)_{u\in U},(\succeq_w, q_w, \ell_w, \mathsf{u}_w)_{w\in W})$ be an instance of SMTI-DIVERSE. We introduce an extra type possessed by each student, and require each college $w \in W$ to have no more than q_w students for this extra type to encode capacities by types.

Motivated by Proposition 4, we will exhaustively branch, for each college and each type vector $\tau \in \{0,1\}^t$, on the choice of a student $\mathsf{wst}(w_j,\tau) \in \mathsf{A}(w_j) \cup \{\top\}$ who will be in $\mathsf{worst}(M,w_j,\tau)$ for a hypothetical feasible and stable matching M. Here $\mathsf{wst}(w_j,\tau) = \top$ is interpreted as $\mathsf{worst}(M,w_j,\tau) = \emptyset$. Moreover we branch to determine the number $\#(w_j,z) \in \{\ell_{w_j}[z],\ldots,\mathsf{u}_{w_j}[z]\}$ of students of each type $z \in [t]$ that each college $w_j \in W$ receives under M.

For each such branch we iteratively try to extend $M_0 = \{\{\mathsf{wst}(w_j,\tau),w_j\} \mid w_j \in W, \tau \in \{0,1\}^t\}$ to a feasible and stable matching which conforms to the guesses in the branch, one not yet matched student at a time.

More specifically, we only add a student-college pair to the matching if doing so maintains the status that each guessed $\mathsf{wst}(w_j,\tau)$-student is least preferred among the students assigned to w_j with type vector τ, the guessed number of students for each college and type is not exceeded, and there is no induced blocking pair involving the added student and some guessed $\mathsf{wst}(w_j,\tau)$-students (as witness). To check these conditions and more importantly to upper-bound the number of considered matchings we keep a *record* in addition to each constructed (partial) matching, the guessed least preferred students $\mathsf{wst}(w_j,\tau)$, $\tau \in \{0,1\}^t$ and the guessed numbers $\#(w_j,z)$ of students, $z \in [t]$, $w_j \in W$. A record for a set U_i of students is an $(m+1)\times(t+1)$-dimensional integer matrix $\mathsf{Q}{\in}\{0,\ldots,n\}^{(m+1)\times(t+1)}$ storing the type-specific number of students assigned to a college, and the number of students assigned to it in total. Two "partial" matchings in a branch can be argued to be equivalent in terms of existence of feasible and stable extensions whenever they have the same record, which is why in each branch we only need to consider at most $n^{(m+1)\cdot(t+1)}$ matchings.

After having considered the last student u_n, we check whether there exists a record Q with a matching M that corresponds to the information in $\#(w_j,z)$, i.e., for each college $w_j \in W$ and each type $z \in [t]$ whether $\mathsf{Q}[j][z] = \#(w_j,z)$ holds. We return M once we found a matching fulfilling the above condition. If no such matching is found, we return that we have a "no"-instance. Correctness can be argued using Proposition 4 and the fact that, in each branch, two partial matchings with the same record are equivalent in terms of existence of feasible and stable extensions. ☐

5 Conclusion

We identified and studied a natural, albeit highly intractable, stable matching problem enhanced with diversity constraints (SMTI-DIVERSE). We showed that while SMTI-DIVERSE is in general Σ_2^{P}-complete, it is polynomial-time solvable when n (the number of students), $m{+}t$ (the number of colleges and types), or $m{+}\mathsf{q}_\infty$ (the number of colleges and the capacity) is a fixed constant.

For future work, studying SMTI-DIVERSE through the lens of parameterized complexity [Downey and Fellows, 2013; Flum and Grohe, 2006; Niedermeier, 2006; Cygan *et al.*, 2015] may provide further insights into the fine-grained complexity of the problem. We left open whether the problem is FPT parameterized by $m+t$. One can also exploit the structure of interactions between students, colleges and types to identify new tractable instances, for instance via the use of notions such as treewidth. Another future research direction is to investigate the trade-off between stability and diversity by allowing few blocking pairs [Abraham *et al.*, 2005; Chen *et al.*, 2018; Mnich and Schlotter, 2020] or few unsatisfied diversity constraints.

Acknowledgments

JC is supported by the WWTF research project (VRG18-012). RG and TH are supported by the Austrian Science Fund (FWF, project P31336).

References

[Abdulkadiroğlu, 2005] Atila Abdulkadiroğlu. College admissions with affirmative action. *International Journal of Game Theory*, 33(535–549), 2005.

[Abraham *et al.*, 2005] David J. Abraham, Péter Biró, and David Manlove. "Almost stable" matchings in the roommates problem. In *Proceedings of the Third International Workshop on Approximation and Online Algorithms (WAOA '05)*, pages 1–14, 2005.

[Ahmed *et al.*, 2017] Faez Ahmed, John P. Dickerson, and Mark Fuge. Diverse weighted bipartite b-matching. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI '17)*, pages 35–41, 2017.

[Aziz *et al.*, 2019] Haris Aziz, Serge Gaspers, Zhaohong Sun, and Toby Walsh. From matching with diversity constraints to matching with regional quotas. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '19)*, pages 377–385, 2019.

[Benabbou *et al.*, 2019] Nawal Benabbou, Mithun Chakraborty, and Yair Zick. Fairness and diversity in public resource allocation problems. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering IEEE*, 42(3):64–75, 2019.

[Berman *et al.*, 2003] Piotr Berman, Marek Karpinski, and Alex D. Scott. Approximation hardness of short symmetric instances of MAX-3SAT. Technical Report 049, ECCC, 2003.

[Bredereck *et al.*, 2014] Robert Bredereck, Jiehua Chen, Piotr Faliszewski, Jiong Guo, Rolf Niedermeier, and Gerhard J. Woeginger. Parameterized algorithmics for computational social choice: Nine research challenges. *Tsinghua Science and Technology*, 19(4):358–373, 2014.

[Chen *et al.*, 2018] Jiehua Chen, Danny Hermelin, Manuel Sorge, and Harel Yedidsion. How hard is it to satisfy (almost) all roommates? In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP '18)*, pages 35:1–35:15, 2018.

[Cygan *et al.*, 2015] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.

[Downey and Fellows, 2013] Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Springer, 2013.

[Flum and Grohe, 2006] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Springer, 2006.

[Garey and Johnson, 1979] Michael R. Garey and David S. Johnson. *Computers and Intractability—A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.

[Gonczarowski *et al.*, 2019] Yannai A. Gonczarowski, Noam Nisan, Lior Kovalio, and Assaf Romm. Matching for the Israeli: Handling rich diversity requirements. In *Proceedings of the 20th ACM Conference on Economics and Computation (ACM EC '19)*, page 321, 2019.

[Hamada *et al.*, 2016] Koki Hamada, Kazuo Iwama, and Shuichi Miyazaki. The Hospitals/Residents problem with lower quotas. *Algorithmica*, 74(1):440–465, 2016.

[Heo, 2019] Eun Jeong Heo. Equity and diversity in college admissions. In *The Future of Economic Design*. Springer, 2019.

[Huang, 2010] Chien-Chung Huang. Classified stable matching. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '10)*, pages 1235–1253, 2010.

[Ismaili *et al.*, 2019] Anisse Ismaili, Naoto Hamada, Yuzhe Zhang, Takamasa Suzuki, and Makoto Yokoo. Weighted matching markets with budget constraints. *Journal of Artificial Intelligence Research*, 65:393–421, 2019.

[Kamada and Kojima, 2015] Yuichiro Kamada and Fuhito Kojima. Efficient matching under distributional constraints: Theory and applications. *American Economic Review*, 105(1):67–99, 2015.

[Kojima, 2019] Fuhito Kojima. New directions of study in matching with constraints. In *The Future of Economic Design*. Springer, 2019.

[Kurata *et al.*, 2017] Ryoji Kurata, Naoto Hamada, Atsushi Iwasaki, and Makoto Yokoo. Controlled school choice with soft bounds and overlapping types. *Journal of Artificial Intelligence Research*, 58:153–184, 2017.

[Laslier *et al.*, 2019] Jean-François Laslier, Hervé Moulin, M. Remzi Sanver, and William S. Zwicker. *The Future of Economic Design*. Springer, 2019.

[Manlove, 2013] David F. Manlove. *Algorithmics of Matching Under Preferences*, volume 2. WorldScientific, 2013.

[Mnich and Schlotter, 2020] Matthias Mnich and Ildikó Schlotter. Stable marriage with covering constraints-A complete computational trichotomy. *Algorithmica*, 82:1136–1188, 2020.

[Nguyen and Vohra, 2019] Thanh Nguyen and Rakesh Vohra. Stable matching with proportionality constraints. *Operations Research*, 67(6):1503–1519, 2019.

[Niedermeier, 2006] Rolf Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006.

[Papadimitriou, 1994] Christos H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[Stockmeyer, 1976] Larry J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1976.