

A scenario-based parametric analysis of the army personnel-to-assignment matching problem

Parametric
analysis of the
army

89

Matthew D. Ferguson,
U.S. Army Human Resources Command, Fort Knox, Kentucky, USA, and
Raymond Hill and Brian Lunday
*Department of Operational Sciences, Air Force Institute of Technology,
Wright-Patterson AFB, Ohio, USA*

Received 22 August 2019
Revised 3 January 2020
26 February 2020
Accepted 27 February 2020

Abstract

Purpose – This study aims to compare linear programming and stable marriage approaches to the personnel assignment problem under conditions of uncertainty. Robust solutions should exhibit reduced variability of solutions in the presence of one or more additional constraints or problem perturbations added to some baseline problems.

Design/methodology/approach – Several variations of each approach are compared with respect to solution speed, solution quality as measured by officer-to-assignment preferences and solution robustness as measured by the number of assignment changes required after inducing a set of representative perturbations or constraints to an assignment instance. These side constraints represent the realistic assignment categorical priorities and limitations encountered by army assignment managers who solve this problem semiannually, and thus the synthetic instances considered herein emulate typical problem instances.

Findings – The results provide insight regarding the trade-offs between traditional optimization and heuristic-based solution approaches.

Originality/value – The results indicate the viability of using the stable marriage algorithm for talent management via the talent marketplace currently used by both the U.S. Army and U.S. Air Force for personnel assignments.

Keywords Assignment problem, Stable marriage algorithm, Solution robustness

Paper type Research paper

1. Introduction

At least twice each year, a group of assignment managers (i.e. human resource specialists) from the U.S. Army's Human Resources Command solve an army personnel-to-assignment matching problem (APAMP), a problem wherein each assignment manager aligned with a career specialty must pair officers with personnel assignments. Each assignment manager may consider the suitability of an officer to fill an assignment, an officer's preference for a

© In accordance with section 105 of the US Copyright Act, this work has been produced by a US government employee and shall be considered a public domain work, as copyright protection is not available. Published in *Journal of Defense Analytics*. Published by Emerald Publishing Limited.

Disclaimer: The views expressed in this article are those of the authors and do not reflect the official policy or position of the U.S. Army, the U.S. Air Force, the Department of Defense or the U.S. Government.

The authors thank the editor and two anonymous referees for their insightful comments and suggestions that improved this manuscript.



given assignment and the alignment of an officer's developmental needs *vis-à-vis* specialized subsets of the assignments related to career development. Optimizing a weighted aggregation of these measures will yield a high *quality* solution to APAMP instances.

Such a problem can be modeled as a classic assignment problem, and the direct optimization of the corresponding mathematical programming formulation will yield an optimal solution based on the preferences and other inputs defined in the problem formulation. However, in the presence of uncertainty regarding the problem structure, the addition and/or removal of elements from either input set will require resolving the assignment problem. Such disruptions occur frequently within APAMP instantiations, as officers may be removed from consideration – or newly considered – for an assignment because of personnel actions (e.g. curtailment of an existing assignment, extension in an existing assignment and retirement) or unit needs (e.g. operation deployment and stop-loss); assignments may become vacant (or filled) on short notice; or a specific officer-to-assignment matching may be directed to meet the needs of a specific unit or senior decision-maker. These disruptions are referred to as perturbations or constraints on the original assignment problem.

The solution to a perturbed APAMP instance may differ notably from the initial solution, and a sequence of such perturbations may induce a level of chaos in the assignment process (Hill *et al.*, 2003), wherein each change in the officer-to-assignment pairings results in a new set of orders, an elevation of officer (and family) stress and a reduced level of officer trust in the assignment process. Usual methods of solving these assignment problems, such as linear programming, are notorious for often drastically changing solutions when problem inputs change in even minor ways.

Thus, when re-solving the APAMP to address a change in the underlying problem structure, we not only want high-quality solutions but also solutions that minimize the number of changes to assignments resulting from instance perturbations, as compared to the initial solution. We define the *robustness* of a solution approach as a measure of how many changes to a solution occur when inputs to the problem change. As the objectives of model robustness and solution quality may conflict, we examine Pareto optimal approaches when considering the tradeoffs between these two objectives.

This study compares classic linear programming solutions to the Gale–Shapley (GS) stable marriage algorithm (SMA) solutions as applied to instances of the APAMP. In particular, the research investigates the tradeoffs between degradation of solution quality and improvement in solution robustness when using an SMA approach for the APAMP. This type of comparison is lacking in current research but is actually a critical component to consider when implementing an SMA for assignment processing. The LP uses an objective function based on preference values, the SMA uses a stability concept, but neither approach explicitly enforces robustness.

This study supports ongoing implementation efforts within both the air force and the army. The Army Talent Management Process recently completed, and many officers received preferred assignments (Rempfer, 2019). Assignment preference is a characteristic important to the army (Army Talent Management Task Force, Headquarters, Deputy Chief of Staff, Army G-1, 2019) that we examine in this study with the intent to provide results that may inform future decisions about the adoption of algorithmic processes to improve assignment outcomes. The Air Force is currently using SMA to support their assignment process and is investigating issues related to preference matrix sparsity, something we look to examine in a sequel to this work; this study focuses on the robustness of assignment solutions using SMA for APAMP.

The remainder of this paper is organized as follows. Section 2 provides an introduction to the matching problem, evaluation metrics that are typically used to compare solutions and variants of the SMA with which solutions may be identified. Section 3 presents necessary notation and sets forth the two categorical approaches used to solve the APAMP: direct optimization of a linear program and the application of the SMA. Section 4 provides and analyzes the results of a set of variants of the two solution methods over a representative set of synthetic test instances, as they relate to solution quality and required computational effort. Section 5 presents conclusions regarding the considered solution methods and proposes directions for further research.

2. Background

The APAMP is a matching problem; it matches officers to assignments to maximize some criterion governing the quality of the matching. For readers familiar with mathematical programming, the APAMP might also be viewed as an assignment problem.

Matching problems arise frequently. [Ahuja *et al.* \(1993\)](#) describe the classical bipartite matching problem as a special instance of the assignment problem wherein the objects are partitioned into two mutually exclusive groups or sets. The two sets of interest are sometimes called *men* and *women*, as is common in an SMA, noted here as sets M and W , respectively, each with complete preferences over the members of the other set. These preferences are used to evaluate the quality of a matching solution. [Bazaraa *et al.* \(2010\)](#) develop a mathematical form of the balanced assignment problem:

$$\begin{aligned}
 & \text{minimize} && \sum_{i=1}^m \sum_{j=1}^m v_{ij} x_{ij} \\
 & \text{subject to} && \sum_{j=1}^m x_{ij} = 1, \quad i = 1, \dots, m, \\
 & && \sum_{i=1}^m x_{ij} = 1, \quad j = 1, \dots, m, \\
 & && x_{ij} \in \{0, 1\}, \quad i, j = 1, \dots, m
 \end{aligned} \tag{1}$$

wherein v_{ij} is the preference (or cost) value associated with the *marriage* or pairing of man i to woman j , and x_{ij} is a binary decision variable associated with the inclusion (i.e. $x_{ij} = 1$) or exclusion (i.e. $x_{ij} = 0$) of that *marriage* from a *matching*. Note that v_{ij} need not equal v_{ji} , and preferences need not be symmetric. Unequal cardinalities of sets M and W are typically handled via the inclusion of placeholder indices or dummy entries that represent a null partner with an appropriate null-valued assignment contribution to the objective.

Another approach to evaluating the relative quality of matching is to use a list of ordinal preferences. [Dean and Swar \(2009\)](#) introduced the *generalized stable allocation problem*, which uses ordinal preferences in allocation problems, but this area of research is relatively unexplored in the literature. An open research question is the value of such ordinal preferences compared to preferences defined on a ratio scale. Such a comparison may be useful when considering the use of ordinal preferences in systems such as the APAMP.

Some literature ascribes value to other measures of matching rather than just fitness derived from costs or preferences. [Kimbrough and Kuo \(2010\)](#) described measures of *equity* and *social welfare*. These measures assign additional value to matchings that are somehow more “fair” to members of each gender. They define equity as the sum of the absolute

differences between the ordinal preferences of each match, and social welfare as the sum of the ordinal preferences of each match.

To illustrate the difference between *social welfare* and *equity*, suppose there are two men, $\{m_1$ and $m_2\}$, and two women, $\{w_1$ and $w_2\}$, with both men and women uniformly preferring indices 1 to 2, with ordinal preference values of 1 and 2, respectively. Observable in Table 1 is that, whereas the two feasible simple matchings have similar social welfare scores, their respective levels of equity differ. Although such nomenclature connotes specific meaning for this example, care should be exercised to determine the applicability of these measures and their correct physical interpretation for a given application and problem instance.

A more commonly examined matching attribute is the notion of *stability*. A matching S is deemed unstable if and only if $\exists (m, w) \notin S$, such that both m and w prefer each other to their respective partners within S (Teo *et al.*, 2001). Such a pair (m, w) is called a *blocking pair* (Dean and Swar, 2009). A stable solution admits no blocking pairs. In Table 1, we observe that, whereas the first matching is stable, the second matching is not. With a view toward solving a matching problem, for any set of ordinal male and female preferences, there exists a stable matching that can be found in quadratic time (Ahuja *et al.*, 1993). Although the literature on exact thresholds to define the related concept of “nearly stable” solutions is sparse (Aldershof and Carducci, 1999), the *relative* stability of S is measured by its total number of blocking pairs (Kimbrough and Kuo, 2010).

The literature on matching and assignment problems is quite rich, covering a wide range of formulations and solution methods. To align with the current research we have focused our review to highlight an important but sparingly covered research area: matchings under preference and stability and the inherent robustness of the returned solution.

The SMA, alternatively known as either the GS or deferred acceptance algorithm, is a heuristic with the objective of finding a stable solution to a bipartite assignment problem. The SMA is well-studied in the literature (Firat *et al.*, 2014). Gale and Shapley (1962) originally proposed it to address the college admissions process, a matching market with incomplete information, to reduce uncertainty for both colleges and applicants. Gale and Shapley reduced the problem to examining potential marriages between an equal number of men and women, each having ordinal preferences over members in the opposite set. For reference, the GS approach is presented as Algorithm 1.

The SMA requires as input two matrices, each representing the ordinal ranking of members of one set on the other (i.e. men on women and vice versa) (Ahuja *et al.*, 1993). These are represented in Algorithm 1 as pref_m and pref_w , respectively.

Of note, the solution approach to determine ordinal preference rankings affects algorithmic performance. Faced with multiple choices (because of equal preferences) the algorithm can choose lexicographical or randomly among the competing choices within a given list. Each approach is examined in this study. Regardless of this decision, either implementation of the GS algorithm yields a stable solution (Ahuja *et al.*, 1993).

Table 1.
Comparison of social welfare, equity and stability

Set of matchings	Social welfare	Equity	Stability
$\{(m_1, w_1), (m_2, w_2)\}$	$2 + 4 = 6$	$0 + 0 = 0$	Yes
$\{(m_1, w_2), (m_2, w_1)\}$	$3 + 3 = 6$	$1 + 1 = 2$	No

Algorithm 1 Pseudocode for the Gale-Shapley Stable Marriage Algorithm

M is the set of men, W is the set of women
 For $m \in M$, $w \in W$: pref_m , pref_w are preference lists for each man, woman
 $S = \emptyset$; the pairs (m, w) in the current matching

```

procedure STABLE MARRIAGE ( $\text{pref}_m$ ,  $\text{pref}_w$ )
    while  $\exists m \in M$  such that  $(m, w) \notin S$  for some  $w \in W$  do
         $m \leftarrow$  select  $m$  s. t.  $m \notin p, \forall p \in S$ 
         $w \leftarrow$  pop ( $\text{pref}_m$ )
         $p \leftarrow p \in S$  where  $w \in p$ 
        if  $w \notin p$  for some pair  $p \in S$  then
             $S \leftarrow (m, w)$ 
        else if  $w$  prefers  $m$  to current match then
            remove pair  $p$  from  $S$ 
             $S \leftarrow (m, w)$ 
        else
            do nothing
    
```

By comparison, [McVitie and Wilson \(1971\)](#) developed a recursive algorithm that provides identical results to the GS algorithm. For test problems of size $m = 50$, they demonstrated a 30% solution time improvement of the recursive algorithm over the GS algorithm while using a programming language favorable to recursion. The [McVitie and Wilson \(1971\)](#) approach also enumerates all stable solutions rather than just the single solution *optimal* to the proposing set (e.g. *male optimal*).

When preferences contain some level of indifference (i.e. when one person has the same utility for a match with multiple partners), the algorithm can accommodate an extended definition of a *blocking pair*. [Manlove \(2002\)](#) explored two variants of the stable marriage problem with ties. A weakly stable matching occurs when there are no blocking pairs that are strictly preferred, whereas a strongly stable matching admits no weakly preferred blocking pairs. Such a situation is easily conceivable within the context of the APAMP.

Perhaps the best known application of the SMA is its use by the National Resident Matching Program (NRMP). The NRMP implements the algorithm to provide stable matchings of medical residents to hospital residency programs and assurances to each program that their slots are filled, without requiring medical students to respond in an unreasonably short time ([Roth, 1984](#)). The success of this particular application earned Roth and Shapley the 2012 Nobel Prize in Economics ([National Resident Matching Program, 2013](#)). Similar uses of deferred acceptance algorithms are noted for several centralized labor markets in the medical community ([Roth, 2008](#)).

Leveraging the literature on matching markets, the Ministry of Education for Singapore uses the SMA to match students to secondary schools using only student preferences and scores on the primary school leaving examination ([Teo et al., 2001](#)). The distribution of students among New York City high schools turned from lottery to a deferred acceptance algorithm in 2003 ([Tullis, 2014](#)). In 2005, Boston public schools adopted a deferred acceptance algorithm to assign students to schools based upon parent preferences while balancing individual school requirements ([Abdulkadirolu et al., 2005](#)).

Applications of the SMA may also be found in modern information technology. A generalized version of the GS algorithm provides a rapid method to establish a matching between users and servers for content delivery networks (e.g. Akamai and Netflix). In content delivery networks, there is very rapid change of preference orderings and

constraints because of network congestion and network topology, both of which are time-variant. The algorithm provides a matching of geographic clusters of users to network nodes that function as content servers that is stable over a short time horizon (Maggs and Sitaraman, 2015).

The SMA has recently received attention for its potential military applications as well. Hill *et al.* (2003) applied the SMA to the nuclear posture review, assigning a notional nonhomogeneous supply of nuclear weapons to a list of targets. Unlike other applications, they observed chaotic behavior of LP-based solutions in dynamic-systems not observed in stable solutions, noting that changes may not be desirable when “adopting a plan already under execution” (Hill *et al.*, 2003). More recently, Naeem and Masood (2010) of the Pakistani Military College of Signals explored the use of the SMA to model detected threats and assign to them a probable target among a set of defended assets. This model enabled a larger threat evaluation and weapon assignment process by quickly providing a robust estimate of likely threat behavior.

Within a human resource context, there have been cases of LP-based approaches in the assignment processes used within the services. In the U.S. Air Force, Wylie (2007) sought to optimize the assignment of so-called *rated officers*, those with aeronautical credentials. The author used an assignment formulation that could be solved as a network flow problem. Recently, Lepird (2016) proposed further research to adapt algorithmic approaches for identifying air force personnel assignments, using either an LP formulation or the SMA. Within the U.S. Army, Sonmez and Switzer (2013) used an agent-based model with an officer-optimal stable solution, wherein side contracts could be chosen to receive a first choice assignment. Although the authors primarily investigated the matching-with-contracts paradigm, the underlying model for their agents used a deferred acceptance algorithm.

The use of the SMA for personnel assignment problems such as the APAMP is relatively underdeveloped. The army currently considers officer and assignment preferences, whereas the air force is more directly investigating the use of SMA to inform the assignment process. The concept of a stable match does provide a nice “cost function” for the assignment process, as both the officer and the duty position are satisfactorily filled. However, there is still much to learn about the subtleties of using the SMA for the military assignment process and how to incorporate various factors into preference functions.

Herein, our study compares the LP and SMA approaches with respect to both the quality and the robustness of attained solutions, an important examination not previously considered in the published literature, but of tremendous utility in the officer assignment process.

3. Models and general solution methods

This section provides the APAMP formulation under consideration. As an additional complicating factor, the APAMP partitions assignments into “key and developmental” (KD) assignments and “broadening” (B) assignments. KD assignments are deemed important to the professional development of an officer, and each officer should have at least one KD assignment during their time at each military rank within a professional career. Accordingly, we define the following sets, parameters and decision variables prior to presenting our model formulation:

Sets

- O = Set of officers considered for assignment;
- O_{KD} = Subset of officers needing a KD assignment;
- O_B = Subset of officers not needing a KD assignment;
- A = Set of assignments needing officers to fill;

A_{KD} = Subset of KD assignments available to fill;
 A_B = Subset of non-KD, broadening assignments available to fill;
 A_D = Set of “dummy” assignments;
 O_D = Set of “dummy” officers; and
 S = Set of determined assignments, called the Matching.

Decision variables

x_{ij} = Binary variable representing the assignment of officer i to assignment j ;
 b_i = Binary variable representing officer i in a blocking pair; and
 z_{ij} = Binary variable representing the assignment of an officer i to assignment j in a lower level program.

Parameters

$|A|$ = Number of total assignments to fill;
 $|AKD|$ = Number of KD assignments available to fill;
 $|AB|$ = Number of non-KD, broadening assignments to fill;
 s = Number of officers seeking assignments;
 $|OKD|$ = Number of officers needing a KD assignment;
 $|OB|$ = Number of officers already having a KD assignment;
 y_i = Year group of officer i from set O ;
 d_i = Officer i relative year group, $d_i = y_i - \min_{i \in O} \{y_i\}$;
 $dmax$ = Largest relative year group among officers;
 P = Preference value matrix, p_{ij} , of officer i to assignment j ; and
 V = Value or preference, v_{ij} of assignment j to officer i .

We note that $O = O_{KD} \cup O_B$ and $O_{KD} \cap O_B = \emptyset$, while similarly $A = A_{KD} \cup A_B$ and $A_{KD} \cap A_B = \emptyset$. When the matching is unbalanced for a given instance, meaning $|A| \neq s$, then either A_D or O_D is populated to a sufficient level to produce a balanced matching. The practical interpretation is that a dummy officer matching to an assignment represents an unfilled assignment whereas an officer matching to a dummy assignment represents an officer who will not receive a new assignment.

The SMA typically refers to the matched objects as men and women. We emulate this terminology for the APAMP and equate men to set O and women to set A . Thus, we can represent the matching implied by each $x_{ij} = 1$ as $(m_i, w_j) \in S$, the matching of an officer i (i.e. man i or m_i) to an assignment j (i.e. woman j or w_j) within the complete set of assignments, S .

The LP and SMA use ordinal officer-to-assignment preferences provided in matrix P . In practice with the army, officers may specify only those jobs they most prefer and those they least prefer, leaving the remainder unspecified. These unranked assignments are referred to as indifferent choices. To model indifference and completely specify P , an average value is used; this computation is the average of the ordinal values that would have been assigned if a strict preference were required. If starting at the n th assignment, an officer expressed indifference between some $k \in \mathbb{N}$ number of assignments, then the preference value is calculated as given in [equation \(2\)](#):

$$p_{ij} = (n - 1) + \sum_{\xi=1}^k \frac{\xi}{k} = n + \sum_{\xi=1}^{k-1} \frac{\xi}{k} \quad (2)$$

For example, if officer i has a clear preference over the first six positions but is indifferent between the four positions ranked as 7, the preference for each of those four positions is calculated using [equation \(2\)](#) as:

$$p_{ij} = 6 + \frac{1 + 2 + 3 + 4}{4} = 7 + \frac{1 + 2 + 3}{4} = 8.5 \quad (3)$$

The LP and SMA also require individual assignment preferences over the set of officers. To inform this parameterization for APAMP, we define each officer's *year group* as the year of their entry into commissioned service, adjusted for prior experience or for unusual career advancement (e.g. promotion ahead of or behind peers during previous promotion boards). Let y_i be the year group of officer $i \in O$, where $d_i = y_i - \min_{i \in O} \{y_i\}$ is an officer's relative year group among officers in the set O , and $d_{max} = \max_{i \in O} \{d_i\}$ is the largest relative year group among officers in the set O , which typically corresponds to the youngest officers. Then, we define a matrix of individual assignment costs in a matrix V , where v_{ij} is the scaled ordinal preference assignment j has for officer i , and is defined as:

$$v_{ij} \triangleq \begin{cases} s d_i, & i \in O_{KD} \text{ and } j \in A_{KD}, \\ s(d_{max} - d_i), & i \in O_B \text{ and } j \in A_B, \\ (1/10)s(d_{max} - d_i), & j \in A_D, \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

The first condition in [equation \(4\)](#) considers KD assignments, for which it assigns a lower cost (i.e. higher scaled ordinal preference) for older officers (i.e. those having lower d_i -values). The second condition considers B assignments, for which it assigns a lower cost (i.e. higher scaled ordinal preference) for younger officers (i.e. those having higher d_i -values). In the third condition of [equation \(4\)](#), null assignments prefer younger officers, albeit with a cost that is scaled by an order of magnitude relative to the B assignment costs. Finally, a cost of 0 is assigned to all other pairings. For all v_{ij} -value computations, the value s is used so that, when considering both officer and assignment preferences in an additive nature, the small range of preferences of assignments over officers is not dominated by the larger number of preferences expressed by officers. From the corporate perspective, the army prioritizes filling KD and then B positions, respectively, with officers who should be aligned within each category, followed by *not* assigning the youngest officers to positions in favor of assigning other officers, because the youngest officers have more time yet to serve in important positions before their next opportunity for promotion. Beyond those preferences, the corporate enterprise is indifferent among ill-suited officer-to-position matchings.

We use four objectives for matching officers to assignments. The first measure matches quality with respect to the usage of KD assignments, the second measures quality with respect to officer preferences, the third measures the stability of a match by giving the number of blocking pairs and the fourth measures quality with respect to assignment preferences:

- *Objective 1: maximize KD opportunity.* [Equation \(5\)](#) prioritizes the full usage of open KD assignments. A desirable matching assigns non-KD complete officers to KD

assignments while minimizing [equation \(5\)](#), which prioritizes the assignment of officers with lower-numbered (i.e. older) year groups:

$$f_1(\mathbf{X}) = \sum_{i \in O_{KD}} \sum_{j \in A_{KD}} y_i x_{ij} \quad (5)$$

- *Objective 2: maximize officer preference satisfaction.* [Equation \(6\)](#) measures the quality of a matching by using the preference information of each officer. Minimizing this quantity seeks the greatest number of officers receiving their highest rated (i.e. lowest ordinal valued) assignments, assuming each officer has a complete preference ordering of assignments:

$$f_2(\mathbf{X}) = \sum_{i \in O} \sum_{j \in A} p_{ij} x_{ij} \quad (6)$$

- *Objective 3: maximize solution stability.* Define the indicator variable b_i for each $i \in O$ as a blocking pair:

$$b_i \triangleq \begin{cases} 1, & (m_i, w_j), (m_k, w_l) \in S, w_l \succ_i w_j, m_i \succ_l m_k \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

Then [equation \(8\)](#) measures the number of blocking pairs present in a solution, where a stable solution has no blocking pairs. To maximize stability, we want to minimize the number of blocking pairs:

$$f_3(\mathbf{X}) = \sum_{i \in O} b_i \quad (8)$$

- *Objective 4: maximize assignment preference satisfaction.* [Equation \(9\)](#) computes the overall assignment preference level for the matching; minimizing [equation \(9\)](#) will maximize satisfaction:

$$f_4(\mathbf{X}) = \sum_{i \in O} \sum_{j \in A} v_{ij} x_{ij} \quad (9)$$

3.1 Solution method #1: linear programming model

Our linear programming model formulation for the APAMP is presented in [equation \(10\)](#):

$$\begin{aligned} & \text{lex min}_{\mathbf{X}} && (f_1(\mathbf{X}), f_2(\mathbf{X}) + f_4(\mathbf{X})) \\ & \text{subject to} && \sum_{i \in O} x_{ij} = 1, && \forall j \in A, \\ & && \sum_{j \in A} x_{ij} = 1, && \forall i \in O, \\ & && x_{ij} \in \{0, 1\}, && \forall i \in O, j \in A. \end{aligned} \quad (10)$$

A solution to [equation \(10\)](#) seeks first to maximize KD opportunities and then, among alternative optimal pairings, to maximize an evenly weighted combination of the satisfaction of officers via their assignment preferences P and the satisfaction of assignments via their scaled officer preferences V . This secondary objective is similar to the social welfare measure of [Kimbrough and Kuo \(2010\)](#) in that it seeks to minimize the collective, mutual regret incurred by a feasible pairing of officers to assignments.

We reformulate [equation \(10\)](#) as [equations \(11\)](#) and (12), for which we can sequentially solve [equation \(12\)](#) to identify the best attainable value of KD assignment usage, k^* , and subsequently solve [equation \(11\)](#):

$$\begin{aligned} \min_{\mathbf{x}} \quad & \sum_{i \in O} \sum_{j \in A} (p_{ij} + v_{ij}) x_{ij} \\ \text{subject to} \quad & \sum_{i \in O} x_{ij} = 1, & \forall j \in A, \\ & \sum_{j \in A} x_{ij} = 1 & \forall i \in O, \\ & \forall i \in O, \\ & x_{ij} \in \{0, 1\}, & \forall i \in O, j \in A. \end{aligned} \tag{11}$$

$$\begin{aligned} \text{where } k^* = \min_{\mathbf{z}} \quad & \sum_{i \in O_{KD}} \sum_{j \in A_{KD}} y_i z_{ij} \\ \text{subject to} \quad & \sum_{i \in O} z_{ij} = 1, & \forall j \in A, \\ & \sum_{j \in A} z_{ij} = 1, & \forall i \in O, \\ & z_{ij} \in \{0, 1\}, & \forall i \in O, j \in A. \end{aligned} \tag{12}$$

Within this study, the LP is modeled using the *numpy* Python programming language package developed by [van der Walt et al. \(2011\)](#) and solved using the *gurobipy* package from Gurobi Optimization (Gurobi Optimization, 2016b), invoking the commercial solver Gurobi v7.0.1 for mixed integer programs (Gurobi Optimization, 2016a).

3.2 Solution method #2: the stable marriage algorithm

For SMA solutions, we consider the following formulation. The SMA is a heuristic and thus does not need a specified formulation of the optimization problem. However, for completeness of presentation, the following formulation represents the SMA approach:

$$\begin{aligned} \text{lex min}_{\mathbf{x}} \quad & (f_1(\mathbf{X}), f_3(\mathbf{X})) \\ \text{subject to} \quad & \sum_{i \in O} x_{ij} = 1, & \forall j \in A, \\ & \sum_{j \in A} x_{ij} = 1, & \forall i \in O, \\ & x_{ij} \in \{0, 1\}, & \forall i \in O, j \in A. \end{aligned} \tag{13}$$

The SMA uses matrix P for officer preferences. The assignment preference matrix, not the matrix C , is constructed considering:

- officers needing KD assignment are preferred for a KD assignment; and
- among those needing KD assignments, officers in the later year group are preferred.

Anyone familiar with military assignments will note that there are many other aspects that could influence an assignment preference: officer skill level, past performance, specific experiences, time on current station, etc. For this study, these were not considered.

3.3 Test design

The data used in the testing was provided by the U.S. Army's Human Resources Command, and it represents preferences for 161 officers over 139 assignments. Accordingly, we create 22 dummy assignments to balance the problem.

Table 2 lists the various solution approach variants examined to subsequently solve the original, or base, matching problem and a problem variant created by augmenting the original problem with randomly generated constraints intended to mimic changes that may actually occur during the assignment process. *LP* and *SMA*, respectively, refer to the use of either a linear program or the SMA as a baseline solution method. *Cold Start* begins the solution process for the perturbed problem without a solution, whereas *Warm Start* initializes the procedure with the incumbent solution from the base problem. As an additional descriptor of the methodology variants, the term *Randomized* indicates that equally preferred assignments in the data set are assigned distinct ordinal preferences randomly, whereas *Lexicographic* breaks such ties by the ordering of assignments within the list. When the algorithm is randomized, it may also be applied repeatedly with the results compared, taking the best result of some number of iterations. This result is determined either by the matching with the fewest number of changes to the incumbent solution, "Fewest Changes," or by the iteration with the lowest objective function value, "Lowest Objective Value."

The base problem is modified with the addition of problem perturbations or constraints. For ten iterations at each solution variant approach listed in Table 2, zero to five random instances of each additional constraint are generated. The incumbent solutions are then compared with the final solutions from the updated model and the number of differences between the solutions is counted to compare the robustness of the SMA and the LP.

Development of realistic assignment scenarios for uncertainties faced by army assignment managers led to the five types of additional constraints under review. While selection of each scenario was experimentally controlled, the selection of the affected officers and assignments was randomized. Each perturbation was generated using the sample function of Python 2.7's random library for the affected officer and/or assignment. This implementation technique provided sampling without replacement with uniform probability. The five types of possible perturbations typically encountered within an army personnel assignment cycle are:

- *Assignment restriction*: An assignment restriction occurs when a subset of A becomes infeasible for an officer. In practice, such an occurrence happens when life events cause an officer to enroll in a program such as the Exceptional Family Member Program or the Married Army Couples Program. These programs may limit the assignment of an officer, even to a desired posting. Another possibility might be duty limitations because of injury or illness preventing assignment to some units, such as airborne status or those deploying in support of overseas operations. A restriction of this type is manifest via the imposition of $x_{ij} = 0$ for officer i and each job, j , deemed infeasible.
- *Directed assignment*: A directed assignment occurs when a particular officer is directed to fill a particular assignment. Whether because of a by-name request from a particular unit or the intervention of a senior decision-maker into the process, these assignments take into account information that may not be routinely available to assignment managers. This type of restriction is imposed by setting $x_{ij} = 1$ for that officer i and specific job j .

- *Rejected match:* Because of how officers are assigned to joint and interservice assignments, also called joint and nominative assignment, some units may exert influence over which officer they receive by effectively exercising a “veto” power. After an assignment manager “nominates” an officer, a unit with such power undertakes its own review or interview process before approving or rejecting the nomination. A restriction of this type requires affixing of the appropriate x_{ij} to values of 0 or 1, respectively, to veto or accept the job j .
- *Unforecasted assignment:* New assignment requirements sometimes occur. This perturbation might occur because of unforeseen organizational changes, such as the establishment of a new joint task force headquarters (e.g. combat operations). Life events of an officer in a high-priority unit, such as a new medical condition or legal action, might also necessitate a reassignment to replace that officer in the current assignment cycle. Such a perturbation increases $|A|$, or $A \leftarrow A \cup \{j_{\text{new}}\}$.
- *Removal of an officer:* Life events, such as medical concerns, legal trouble or personal career decisions might impact the availability of an officer for reassignment. Despite many media portrayals, the army’s processes account for these in its reassignment decisions. Officers might retire in-lieu of a reassignment or request stabilization so that a high school-aged child does not have to move to a new school for their senior year. This perturbation entails either removing officer i from set O (i.e. $O \leftarrow O \setminus \{i\}$) or setting $x_{ij} = 0$ for all jobs j for that specific officer i .

Three response measures are used: solution time, solution quality in both problems and the number of changes between the solution to the original problem and the solution of the perturbed problem.

4. Testing, results and analysis

The first comparison focuses on the computational effort required by each of the solution approaches considered. Table 3 summarizes the computational effort for each solution method as measured by the elapsed time on an Intel® Core™ i5-7200U CPU (2 cores) with hyperthreading enabled running Ubuntu 16.10 in an Anaconda 4.2-managed Python 2.7.12 environment. Not surprisingly, the SMA heuristics run the quickest while the SMA with the randomized selections runs the longest. From a practical perspective, all the techniques run sufficiently quickly – in seconds – indicating that we can focus the comparison on the solution quality and solution robustness characteristics.

Solution quality is the second metric of interest. Table 4 summarizes the results, providing the solution methods, ordered by average objective function value and with an

Table 2.
Variants of SMA and LP under test

Variant	Abbreviation
LP, cold start	LP-Cold
LP, warm start	LP-Warm
SMA, randomized with cold start and fewest changes of 10 iterations	SMA-10r
SMA, randomized with cold start and fewest changes of 30 iterations	SMA-30r
SMA, randomized with cold start and lowest objective value of 5 iterations	SMA-5o
SMA, randomized cold start and fewest changes of 5 iterations	SMA-5r
SMA, lexicographic with cold start	SMA-Lex
SMA, lexicographic with warm start	SMA-Lex, Warm
SMA, randomized with warm start	SMA-Warm

indication of significant differences based on a Scheffé multiple comparisons method. The LP approaches clearly provide solutions with the best objective function values, and these values are significantly better than any of the SMA-based approaches. However, the SMA solutions are still within 2.5% of the LP solutions.

As it pertains to the use of LP-based approaches, Table 5 examines whether using a warm start helps with respect to solution quality. The results in Table 3 depict no statistically significant difference in computation times, and the results in Table 5 likewise indicate no significant difference in the resulting solution quality. With LP-based methods, the average officer preference in the matching S indicates officers can expect their 16th assignment choice. This behavior is pathological and is not uncommon to LP assignment approaches; the LP is indifferent to the individual and focused on the overall metric. Naturally, a more common metric of matching quality among assignment managers might be the percentage of officers receiving one of their top-three choice assignments, as shown in Figure 1. We note that all of the SMA algorithms outperform the LP mode with respect to this informal metric. However subsequent LP models could be tuned to narrow the gap. In short, while LP models outperform when using the established objective function, SMA models appear to consistently provide a greater percentage of the population with desired assignments.

The final metric of comparison is the robustness in the solution method; i.e. how much does the solution change, given perturbations to the underlying problem? Figure 2 illustrates the mean number of changes to an incumbent solution for a number of randomly

Solution method	95% Scheffé CI for mean times (s)
SMA, fewest changes of 30 iterations (SMA-30c)	3.0069 ± 0.121
SMA, fewest changes of 10 iterations (SMA-10c)	1.1289 ± 0.034
SMA, best objective of 5 iterations (SMA-5o)	0.6592 ± 0.021
SMA, fewest changes of 5 iterations (SMA-5c)	0.6592 ± 0.022
LP, warm start (LP-W)	0.6453 ± 0.020
LP, cold start (LP-C)	0.6404 ± 0.022
SMA, lexicographic (SMA-Lex)	0.2811 ± 0.009
SMA, randomized with warm start (SMA-Warm)	0.2744 ± 0.006
SMA, lexicographic with warm start (SMA-Lex, Warm)	0.2682 ± 0.014

Table 3.
Computational effort
as elapsed time with
95% Scheffe CIs

Treatment	Mean	Objective function value	
		95% CI	group*
SMA-Lex	72,780		<i>a</i>
SMA-Lex, Warm	72,460		<i>b</i>
SMA-5o	72,400		<i>bc</i>
SMA-5c	72,380		<i>c</i>
SMA-10c	72,380		<i>c</i>
SMA-30c	72,360		<i>c</i>
SMA-Warm	72,240		<i>d</i>
LP-C	70,980		<i>e</i>
LP-W	70,960		<i>e</i>

Table 4.
Solution quality with
groupings from 95%
Scheffé CIs

Note: *Treatments sharing a letter are statistically equivalent

chosen perturbations, where each of the three types of perturbations is equally likely to occur. Table 6 provides the corresponding mean number of changes with 95% Scheffé confidence intervals (CIs). From these results, we infer very similar robustness qualities among the LP-based and the randomized SMA approaches but a notable superiority of the lexicographic SMA approach variants. Thus, as was found in Hill *et al.*'s (2003) study, the SMA approach is quite robust, essentially limiting changes to the solution when the original problem is perturbed.

Table 5.
Solution quality LP
methods, with 95%
Scheffé CIs

Variant	Mean objective	Avg. officer preference
LP warm start (LP-W)	70,960 ± 265	16.04 ± 0.77
LP cold start (LP-C)	70,980 ± 301	16.09 ± 0.78

Figure 1.
Percentage of officers
matched to a top-
three preferred
assignment





Figure 2.
Robustness (changes
to incumbent
solution) by method



Based upon the observed data, a trade-off appears to exist between reductions in the number of changes and reductions in objective function value. Lexicographic SMA and warm-start LP are both Pareto-optimal approaches in a bi-objective comparison of quality and robustness. For the two objectives of quality and robustness, an approach is Pareto-optimal if changing approaches to improve one objective necessitates a decrease in the value of the other objective (Ehrgott, 2006). Figure 3 compares the LP-warm, SMA-lex and SMA-warm approaches along the dimensions of solution quality and solution robustness. For each of the three plots, the lower left corner of each plot is the utopia point, optimizing both solution robustness and solution quality. In addition to the comparative depiction of the tradeoff of quality vs robustness for each of the approaches, Figure 3 also illustrates clearly the lower variability in solution quality and robustness associated with the warm-start SMA approach.

Group	Solution method	95% Scheffé CI for mean times (s)
<i>a</i>	SMA, randomized with best objective of 5 iterations	84.40 ± 6.94
<i>b</i>	SMA, randomized with fewest changes of 5 iterations	76.52 ± 5.06
<i>c</i>	SMA, randomized with fewest changes of 10 iterations	73.80 ± 4.73
<i>d</i>	SMA, randomized with fewest changes of 30 iterations	70.51 ± 4.12
<i>e</i>	LP	62.31 ± 8.18
	LP with warm start	61.64 ± 8.24
<i>f</i>	SMA, lexicographic	48.89 ± 8.76
<i>g</i>	SMA, lexicographic with warm start	10.79 ± 5.12
	SMA, randomized with warm start	10.39 ± 4.76

Table 6.
Mean number of changes with 95% Scheffé CIs




Figure 3.
Developing the efficient frontier of solution techniques for mean solution quality and mean robustness

5. Summary and conclusions

Personnel assignment processes are often based solely on the needs of the organization. Recent initiatives attempt to provide greater consideration of personal preferences. The organizations' needs are also of importance, and so an appropriate compromise is inform assignments using both the organizations' and the officers' preferences. To this end, the SMA has been proposed as a potential alternative solution method to the assignment process and, in particular, the APAMP considered in this paper. Adoption of the proposed assignment approach does require deeper insight into what the solution approach provides in comparison to current methods. Within this context, this study compares several variants of the traditional linear programming approach and the SMA solution approach with respect to three performance metrics: solution time, solution quality as measured by an objective function value and solution robustness as measured by the number of changes required between a solution to the APAMP and the solution to perturbed variant thereof.

Whereas none of the solution variants examined are challenged with respect to required solution time, the results from this study demonstrate that a tradeoff exists between the quality and robustness of solutions when considering an LP vs an SMA solution approach. In testing herein, an LP approach yields the highest average objective function value but requires more changes to a solution when a problem is perturbed. By comparison, an SMA approach using the incumbent solution to the original problem as a warm start attains solutions within 2-3% of the quality corresponding to LP-based solutions while significantly reducing the number of changes required from the original solution upon problem perturbation.

Interestingly, the APAMP is evolving in its application. In the fall of 2017 and for the first time, the Officer Personnel Management Directorate (OPMD) of Human Resources Command required all officers being reassigned in 2018 to indicate a strict ordinal preference over all available assignments. The first use of these officer preferences completed in December 2019 (Rempfer, 2019). OPMD has likewise asked for units to indicate some preference over the set of officers. Such input requirements lend themselves to the direct application of the methods considered herein, which portends a more stable assignment process in the presence of the identified uncertainties.

There are a number of recommendations for extensions to this work. While using ordinal preferences, the preference matrix used should be complete. Thus, a mechanism would be required to allow officers to fully specify their preferences while an algorithm would be required for each assignment to fully specify its preference structure. We demonstrated an algorithm using a subcategory of assignments (i.e. KD assignments) as an additional factor, and future algorithmic examinations should consider additional factors of importance to assignment teams. Alternatively, an examination of a continuous preference scale *vis-à-vis* an ordinal scale has merit. In lieu of preference elicitation to achieve complete preference matrices, modifications to SMA algorithms should be examined to accommodate sparse preference matrices. Whereas traditional applications of SMA have some associated strategy-related proofs (Roth, 1984), further investigation into the game-theoretic aspects of preference elicitation is warranted. Finally, this study also adopted a strict pre-emptive goal programming formulation and generated problem perturbations under an assumption of equal probability; either of these assumptions might be relaxed upon further interviews with personnel assignment subject matter experts to further refine the respective solution approaches and modeling assumptions.

References

- Abdulkadirolo, A., Pathak, P.A., Roth, A.E. and Snmez, T. (2005), "The Boston public school match", *American Economic Review*, Vol. 95 No. 2, pp. 368-371, available at: www.jstor.org/stable/4132849

-
- Ahuja, R., Magnanti, T. and Orlin, J. (1993), *Network Flows: theory, Algorithms, and Applications*, Prentice Hall, Upper Saddle River, NJ.
- Aldershof, B. and Carducci, O.M. (1999), "Stable marriage and genetic algorithms: a fertile union", *Journal of Heuristics*, Vol. 5 No. 1, pp. 29-46.
- Army Talent Management Task Force, Headquarters, Deputy Chief of Staff, Army G-1 (2019), "Army officer talent management", available at: www.army.mil/standto/2019-04-26
- Bazaraa, M., Jarvis, J. and Sherali, H. (2010), *Linear Programming and Network Flows*. Wiley.
- Dean, B.C. and Swar, N. (2009), "The generalized stable allocation problem", *International Workshop on Algorithms and Computation*, Springer, pp. 238-249.
- Ehrgott, M. (2006), *Multicriteria Optimization*, Springer Science and Business Media, Auckland.
- Firat, M., Hurkens, C. and Laugier, A. (2014), "Stable multi-skill workforce assignments", *Annals of Operations Research*, Vol. 213 No. 1, pp. 95-114.
- Gale, D. and Shapley, L. (1962), "College admissions and the stability of marriage", *The American Mathematical Monthly*, Vol. 69 No. 1, pp. 9-15, available at: www.jstor.org/stable/2312726
- Gurobi Optimization (2016a), "Gurobi optimizer example tour", available at: www.gurobi.com/documentation/7.0/examples.pdf (accessed 29 October 2016)
- Gurobi Optimization (2016b), "Gurobi optimizer features and benefits", available at: www.gurobi.com/products/features-benefits (accessed 29 October 2016).
- Hill, R.R., Chopra, V., Nadkarni, B., Cotsworth, W. and Schroeder, G. (2003), "Examining the stable marriage algorithm paradigm for agent-based optimization applications in support of defense planning challenges", IIE Annual Conference. Proceedings, *Institute of Industrial Engineers, Houston, TX*.
- Kimbrough, S.O. and Kuo, A. (2010), "On heuristics for two-sided matching: revisiting the stable marriage problem as a multiobjective problem", *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation, ACM*, pp. 1283-1290.
- Lepird, J. (2016), "Position paper on a next generation air force personnel assignment system", Tech. rep., United States Air Force.
- McVitie, D.G. and Wilson, L.B. (1971), "The stable marriage problem", *Communications of the Acm*, Vol. 14 No. 7, pp. 486-490.
- Maggs, B.M. and Sitaraman, R.K. (2015), "Algorithmic nuggets in content delivery", *ACM SIGCOMM Computer Communication Review*, Vol. 45 No. 3, pp. 52-66.
- Manlove, D.F. (2002), "The structure of stable marriage with indifference", *Discrete Applied Mathematics*, Vol. 122 Nos 1/3, pp. 167-181.
- Naeem, H. and Masood, A. (2010), "An optimal dynamic threat evaluation and weapon scheduling technique", *Knowledge-Based Systems*, Vol. 23 No. 4, pp. 337-342.
- National Resident Matching Program (2013), "The Sveriges riks bank prize in economic sciences in memory of Alfred nobel", available at: www.nrmp.org/wp-content/uploads/2013/08/The-Sveriges-Riksbank-Prize-in-Economic-Sciences-in-Memory-of-Alfred-Nobel1.pdf (accessed 29 October 2016).
- Rempfer, K. (2019), "Almost half of officers match to top job choice under new system", *Army Times*, Vol. 81 No. 1, p. 8.
- Roth, A.E. (1984), "The evolution of the labor market for medical interns and residents: a case study in game theory", *Journal of Political Economy*, Vol. 92 No. 6, pp. 991-1016.
- Roth, A.E. (2008), "Deferred acceptance algorithms: history, theory, practice, and open questions", *International Journal of Game Theory*, Vol. 36 Nos 3/4, pp. 537-569, doi: [10.1007/s00182-008-0117-6](https://doi.org/10.1007/s00182-008-0117-6), available at: <http://dx.doi.org/10.1007/s00182-008-0117-6>
- Sonmez, T. and Switzer, T.B. (2013), "Matching with (branch-of-choice) contracts at the United States military academy", *Econometrica*, Vol. 81 No. 2, pp. 451-488, doi: [10.3982/ECTA10570](https://doi.org/10.3982/ECTA10570). available at: <http://dx.doi.org/10.3982/ECTA10570>

- Teo, C.P., Sethuraman, J. and Tan, W.P. (2001), "Gale-Shapley stable marriage problem revisited: strategic issues and applications", *Management Science*, Vol. 47 No. 9, pp. 1252-1267.
- Tullis, T. (2014), "How game theory helped improve New York City's high school application process", available at: www.nytimes.com/2014/12/07/nyregion/how-game-theory-helped-improve-new-york-city-high-school-application-process.html, accessed: 2016-10-29
- van der Walt, S., Colbert, S.C. and Varoquaux, G. (2011), "The numpy array: a structure for efficient numerical computation", *Computing in Science and Engineering*, Vol. 13 No. 2, pp. 22-30, doi: [10.1109/MCSE.2011.37](https://doi.org/10.1109/MCSE.2011.37).
- Wylie, A.M. (2007), "Optimization of rated officer staff assignments", Master's thesis, Air Force Institute of Technology.

Corresponding author

Brian Lunday is the corresponding author and can be contacted at: brian.lunday@afit.edu