

# MatchMaker: A DSL for Game-Theoretic Matching

Prashant Kumar and Martin Erwig

Oregon State University, Corvallis OR 97330, USA  
{kumarpra,erwig}@oregonstate.edu

**Abstract.** Existing tools for solving game-theoretic matching problems are limited in their expressiveness and are difficult to use. In this paper, we introduce MATCHMAKER, a Haskell-based domain-specific embedded language (DSEL), which supports the direct, high-level representation of matching problems. Haskell’s type system, especially the use of multi-parameter type classes, facilitates the definition of a very general interface to matching problems, which can be quickly instantiated to a wide variety of different matching applications. As another novel contribution, MATCHMAKER provides combinators for dynamically updating and modifying problem representations and for analyzing matching results.

## 1 Introduction

A large class of problems are instances of matching problems. Examples include the assignment of children to different schools, students to universities and campus housing, or doctors to hospitals, kidney transplant patients to donors, and many others. In each of these problems, the participants in the matching process typically have a *preferences* over the entities they are matched to, and the task is to find a matching that is in some sense optimal with respect to these preferences. In addition to the traditional applications of these problems, the proliferation of the internet has made the collection of preferences easier, thereby opening new application domains for matching problems. The importance of matching is also highlighted by the fact that the 2012 Nobel Prize in Economics was awarded to Lloyd S. Shapley and Alvin E. Roth for their work on stable matching problems.

Despite its apparent usefulness, the actual software support for expressing and solving matching problems is surprisingly limited in a number of ways. For example, the currently available software tools for solving matching problems are limited in expressiveness and often difficult to use. Almost all the available matching libraries use strings to encode the matching problem, which affects readability and maintainability of the encoded problems. Employing untyped representations limits the options for checking the validity of the encoding and producing meaningful error messages. As we will demonstrate, MATCHMAKER leverages Haskell’s rich type system and its type class system to facilitate high-level representations of matching problems that are readable, easily modifiable, and provide good error checking.

Arthur	Sunny	Joseph	Latha	Darrius
City	City Mercy	City General Mercy	Mercy City General	City Mercy General

(a) Applicants' hospital preferences

Mercy	City	General
Darrius Joseph	Darrius Arthur Sunny Latha Joseph	Darrius Arthur Joseph Latha

(b) Hospitals' ranking of applicants

Fig. 1: Matching hospitals with applicants: a two-sided stable matching example.

39 MATCHMAKER already implements algorithms for a large class of matching  
40 problems. More specifically we implement *bipartite stable matching with two-*  
41 *sided preferences*, *bipartite stable matching with one-sided preferences*, and *same-*  
42 *set matching problems with one-sided preferences*. Together these represent the  
43 most important and widely applicable matching problems [16, 11, 7]. However,  
44 due to space constraints we only present the encoding of the bipartite stable  
45 matching problem in the paper.

46 Our DSL makes the following main contributions. It:

- 47 – offers high-level, type-safe, extensible representation for matching problems.
- 48 – defines a scalable mechanism for describing preferences based on function  
49 definitions and abstract criteria.
- 50 – provides functions to analyze and compare the results of various matchings.
- 51 – is easily extensible to represent new matching problems.

52 The remainder of this paper is structured as follows. In Section 2 we introduce  
53 stable bipartite matching problems with two-sided preferences and encode them  
54 in MATCHMAKER with explicit preferences. In Section 3 we illustrate how to  
55 represent preferences implicitly using Haskell's abstract data types and func-  
56 tions. In Section 4 we introduce combinators to update the existing matching  
57 representations plus functions for comparing the results of two matchings. In  
58 Section 6, we compare MATCHMAKER to other tools for matching. Finally, in  
59 Section 7 we provide conclusions.

## 60 2 Bipartite Stable Matching With Two-Sided Preferences

61 Consider the problem of assigning applicants to hospitals, taken from NRMP's  
62 website [12] for illustration. Figure 1 shows that each of the three hospitals have a  
63 quota of 2. There are five applicants wishing to get a residency. Hospitals provide  
64 preference for applicants and vice versa. The stable matching algorithm (also  
65 called *delayed acceptance algorithm* [14, 4]) comes up with a match of residents  
66 to hospitals having the following two characteristics:

- 67 (1) Each applicant is assigned to only one hospital and no hospital is assigned  
68 more applicants than its quota.
- 69 (2) The resulting match is *stable*. This condition is described in the delayed  
70 acceptance algorithm as the match not having a *blocking pair*, which is a  
71 pair of hospital and applicant currently assigned to different partners but  
72 who prefer each other more than their current assignment. The presence  
73 of such pairs undermines the effectiveness of the matching process as these  
74 pairs can make private arrangements leaving behind their partners assigned  
75 by the matching algorithm. [16, Chapter 5] shows an example of an unsta-  
76 ble matching mechanism for matching doctors to hospitals in Birmingham  
77 and Newcastle used in 1960s and 1970s. The instability of the outcome lead  
78 to doctors and hospitals entering private negotiations outside the match-  
79 ing process which left many doctors without a position and many hospitals  
80 without a resident. This culminated into abandonment of the mechanism.  
81 Gale and Shapley [4] showed that a special property of bipartite matching  
82 markets is that stable matchings always exist.

83 Let us try to match hospitals with applicants taking into account their prefer-  
84 ences and quotas. Consider the preference list of Darrius. He prefers City the  
85 most, and City also ranks him the highest amongst the candidates. It is easy  
86 to deduce that Darrius will end up at City. Now, if we look at the preference  
87 list of Sunny, we see that she considers just City and Mercy for her residency,  
88 but Mercy doesn't rank her meaning that she can't be assigned there. Her first  
89 option, City, does rank her third. However, notice that the two people ranked  
90 above her, Darrius and Arthur, have listed City as their first choice. If they  
91 are assigned the two positions, then Sunny is left without an offer as Mercy is  
92 the only hospital in her preference set that also ranks her. Could we have ac-  
93 commodated Darrius at Mercy leaving room for Sunny at City? Although this  
94 does lead to Sunny getting accommodated at City, it leads to an instability in  
95 the matching process due to formation of a blocking pair of Darrius and City:  
96 Darrius still prefers City over Mercy and City still prefers Darrius over Sunny,  
97 that is, they profit from forming their own match leaving behind their assigned  
98 matches.

99 In this section, we demonstrate how to represent this example in MATCH-  
100 MAKER and generate stable matching. To motivate the different design choices,  
101 it is instructive to look at the formal model of stable matching.

## 102 2.1 Modeling Stable Matching

A two-sided stable matching problem between applicants and hospitals is a 6 tu-  
ple  $(A, H, P_A, P_H, Q_A, Q_H)$  where  $A = \{a_1, a_2, \dots, a_m\}$  and  $H = \{h_1, h_2, \dots, h_n\}$   
represents the finite disjoint sets of applicants and hospitals, respectively [16].  
The preference of each applicant  $a \in A$  is represented by an ordered list of  
preferences  $P(a)$  on set  $H$ . Similarly, the preference of each hospital  $h \in H$   
is represented by an ordered list of preferences  $P(h)$  on set  $A$ . The set of all

preference list for applicants and hospitals is defined as  $P_A$  and  $P_H$  as shown below.

$$P_A = \{P(a_1), P(a_2), \dots, P(a_m)\} \text{ and } P_H = \{P(h_1), P(h_2), \dots, P(h_n)\}$$

Each hospital  $h$  is also assigned a positive integer  $Q(h)$ , also called its *quota*, that represents the maximum number of applicants it could admit. Similarly, each applicants is also assigned a quota.

$$Q_A = \{Q(a_1), Q(a_2), \dots, Q(a_m)\} \text{ and } Q_H = \{Q(h_1), Q(h_2), \dots, Q(h_n)\}$$

103 For the applicant-hospital matching problem, it is obvious that applicants have a  
104 quota of 1, since they can work at only one hospital. However, in other examples  
105 of matching problems with two-sided preferences, both the sets can have quotas  
106 greater than 1.

107 A *matching* is a relation  $\mu \subseteq A \times H$  that satisfies the following two conditions:  
108 (1)  $\forall a \in A : |\mu(a)| \leq Q_A(a)$  and  $\forall h \in H : |\mu(h)| \leq Q_H(h)$  and (2)  $\mu(a) = h \Leftrightarrow$   
109  $a \in \mu(h)$ . The first condition ensures the matching satisfies the quota restrictions,  
110 and the second condition ensures consistency. In our current example that means  
111 that a hospital is in an applicant's match if and only if that applicant is also in  
112 the hospital's match.

113 The formal model guides the design of our DSL, which we demonstrate with  
114 the help of our example next.

## 115 2.2 DSL Representation of Matching Problems

116 The first step in encoding our example is to represent the two sets to be matched  
117 as Haskell data types.

```
data Applicant = Arthur | Sunny | Joseph | Latha | Darrius
data Hospital = City | Mercy | General
```

118 The definitions and type signatures of various data types, type classes, and  
119 functions used in this section are summarized in Figure 2a. We need a way  
120 to store the preferences of the applicants and hospitals for which we use two  
121 mappings **Rec** (abbreviated for record) and **Info**, which is a collection of **Recs**,  
122 represented as a mapping.

123 Specifically, **Info Applicant Hospital Rank** maps every applicant to a record  
124 **Rec Hospital Rank**, which maps hospitals to their ranks as specified by the  
125 applicant. Similarly, the ranking of applicants by hospitals is represented in  
126 **Info Hospital Applicant Rank** where the individual preferences of each hospital  
127 are recorded in the mapping **Rec Hospital Rank**.

128 The multi-parameter type class **Preference a b c** provides an interface to  
129 specify the preferences. An interesting aspect of the class definition is the func-  
130 tional dependency specification, which signifies that types **a** and **b** uniquely de-  
131 termine type **c**. The **gather** function of the type class captures the preference of  
132 elements of set **a** for elements of set **b** using a type **c** and stores it in the mapping

```

type Capacity = Int

forall :: Capacity -> a -> Capacity
forall c _ = c

class (Bounded a, Enum a, Ord a) => Set a where
  members :: [a]
  members = enumFromTo minBound maxBound

  quota :: a -> Capacity
  quota = forall 1

data Rec b c = Rec {unRec :: M.Map b (Maybe c)}
data Info a b c = Info {unInfo :: M.Map a (Rec b c)}

data Rank = Rank {unRank :: Int}

class Preference a b c | a b -> c where
  gather :: Info a b c

info :: Ord2 a b => [(a, [(b, c)])] -> Info a b c
choices :: Ord2 a b => [(a, [b])] -> Info a b Rank

data Match a b = Match {unMatch :: [(a, [b]), Maybe Capacity]}

ranks :: (Set2 a b, Norm c, Weights a) => Info a b c -> Match
twoWay :: (Preference2 a b c d, Set2 a b, Norm2 c d) => Match

twoWayWithCapacity :: (Preference2 a b c d, Set2 a b, Norm2 c
  Match a b
twoWayWithPref :: (Preference2 a b c d, Set2 a b, Norm2 c
  Info a b c -> Info b a d -> Match a b

class Norm a where
  components :: a -> [Double]
  components _ = []

  norm :: (a, Maybe Double) -> Double
  norm = sum . components . fst

instance Norm Double where
  norm (v, nv) = min (v / (fromJust nv)) 1

instance Norm Rank where
  norm (Rank r, _) = 1 / fromIntegral r

instance Norm Bool where
  norm (x, Nothing) = if x then 1.0 else 0.0
...
outOf :: Norm a => a -> Double -> Double
outOf x y = norm (x, Just y)

only :: Norm a => a -> Double
only x = norm (x, Nothing)

class Weights a where
  weights :: a -> [Double]
  weights _ = [1.0]

class Weights a => Preference a b c | a b -> c where
  gather :: Info a b c

```

(a) Definitions for encoding and storing (b) Support for Representational Rank-preferences.

Fig. 2: Major Definitions of MATCHMAKER.

133 **Info** a b c. The smart constructors **info** and **choices** are used to construct the  
 134 preference mappings from list of tuples.<sup>1</sup>

135 Our current example requires two-sided specification of preferences. This en-  
 136 tails two instance definitions of the **Preference** class, one for specifying hospitals’  
 137 preferences and one for applicants’. However, before presenting those definitions,  
 138 we discuss a particular design choice for the type class. One question is whether  
 139 we should have simplified the definitions of the **Info** mappings and consequently  
 140 the **Preference** class by removing their last type argument and hard-coding the  
 141 **Rank** in the definitions instead. This would mean that the rank of an item is spec-  
 142 ified by the position of that item in a list. While this does simplify the design,  
 143 the constraint to relate the items being matched in just one way also limits the  
 144 expressivity of the domain. The advantages of our design choice becomes appar-  
 145 ent in Section 3 where we instantiate the third argument of **Info** and **Preference**  
 146 with richer types than **Rank** that allow agents to implicitly rank other agents,  
 147 which eases the cognitive burden and effort in coming up with a preference list.

148 The ranked preference lists of applicants for hospitals can be specified with  
 149 an appropriate instance of **Preference** type class using the **choices** function as

<sup>1</sup> We mostly show only the type signatures and present implementations only when they contribute to a better understanding. For the complete code, see <https://github.com/AnonymousForConfReview/MatchMaker>.

150 shown below. The infix operation `-->` is simply syntactic sugar for building pairs.  
 151 In this representation, rankings are based on positions. For example, the fact that  
 152 City precedes Mercy in the preference list of Sunny means that she prefers City  
 153 to Mercy.

```
instance Preference Applicant Hospital Rank where
  gather = choices [Arthur --> [City],
                  Sunny --> [City, Mercy],
                  Joseph --> [City, General, Mercy],
                  Latha --> [Mercy, City, General],
                  Darrius --> [City, Mercy, General]]
```

154 The ranked preference lists of applicants for hospitals can be similarly encoded.

```
instance Preference Hospital Applicant Rank where
  gather = choices [Mercy --> [Darrius, Joseph],
                  City --> [Darrius, Arthur, Sunny, Latha, Joseph],
                  General --> [Darrius, Arthur, Joseph, Latha]]
```

155 Finally, to encode the quota information we define a type class called `Set` with  
 156 `quota` as a member function as shown in Figure 2a. We also define a function  
 157 `members` that can list all the elements of a set. The function `forall` is used to  
 158 assign the same quota to every member of the set to be matched.

159 The instances of `Set` for the `Applicant` and the `Hospital` type are shown below  
 160 where each hospital is assigned a quota of 2 and each applicant is assigned a  
 161 default quota of 1.

```
instance Set Applicant
instance Set Hospital where quota = forall 2
```

## 162 2.3 Generating Stable Matchings

163 In general, a matching problem can have multiple stable matchings. However,  
 164 two are especially significant. For our problem, these are the *hospital-optimal*  
 165 *stable match* and the *applicant-optimal stable match*. (Sometimes the adjective  
 166 “stable” is omitted for brevity.) In a hospital-optimal match, hospitals do as  
 167 good as they possibly can. While not intended, the structure of the matching  
 168 problem entails that a stable match where hospitals perform their best is also  
 169 a stable match where applicants perform their worst [16, Chapter 2, Corollary  
 170 2.14]. Similarly, in an applicant-optimal match applicants perform their best  
 171 and hospitals their worst. Interestingly, the NRMP program was shown to be  
 172 hospital optimal [21] before it was changed to be applicant optimal in 1997 [15].

173 A stable match can be computed with the function `twoWayWithPref`, which  
 174 takes two preference encodings of type `Info` type and yields a value of type  
 175 `Match a b` that stores all the elements of set `b` matched to an element of set `a`. The  
 176 overloaded value `twoWay` triggers the computation by inferring the `Info` arguments  
 177 from its type annotation. For example, the annotation `Match Applicant Hospital`  
 178 generates the applicant-optimal matching.

```
> twoWay :: Match Applicant Hospital
{Sunny --> [],
 Darrius --> [City],
 Latha --> [General],
 Joseph --> [General], Arthur --> [City]}
```

179 Similarly, `Match Hospital Applicant` generates a hospital-optimal matching.

```
> twoWay :: Match Hospital Applicant
{City --> [Arthur,Darrius],
 Mercy --> [],
 General --> [Latha,Joseph]}
```

180 Comparing the two matchings shows that they are the same. However, this need  
181 not always be the case.

182 The DSL also provides a function `twoWayWithCapacity` to find the remaining  
183 quotas in a matching. The next example shows that General and City have  
184 exhausted their quotas of applicants, whereas Mercy's quota of 2 is untouched,  
185 since no residents have been assigned to it.

```
> twoWayWithCapacity :: Match Hospital Applicant
{City --> [Arthur,Darrius] : 0, Mercy --> [] : 2, General --> [Latha,Joseph] : 0}
```

### 186 3 Representational Ranking

187 There is something unsatisfactory about the encoding of the NRMP example  
188 in the last section. Specifically, the fact that every hospital needs to provide a  
189 rank for every applicant and vice versa is problematic. For example, consider a  
190 hospital with 100 applications for its residency positions. How does the admission  
191 committee determine the specific rank for each applicant? Similar problems arise  
192 for applicants applying to a large number of hospitals.

193 Instead of ranking through an ordered list, one can often describe a ranking  
194 through a function that computes a rank based on attributes of the elements  
195 to be ranked. For example, a hospital might specify that it would like to rank  
196 the candidates based on their MCAT scores, their performance in the interview  
197 with the hospital, their previous experiences in the field, and whether or not  
198 their previous degree is from their hospital assigning weights of say 40%, 30%,  
199 20%, and 10%, respectively. A score can be generated for each candidate using  
200 such a formula, and the reciprocal of this score is then used to obtain rank  
201 for a particular candidate. Every hospital may assign different weights for the  
202 different criteria. Some may even decide to not use a particular criterion at all  
203 (which is the same as assigning it a weight of 0). MATCHMAKER facilitates this  
204 form of ranking. To this end, we define a data type `AInfo` for storing the relevant  
205 applicant data.

```
data AInfo = Appl {examScore      :: Double,
                  experience      :: Double,
                  interviewScore  :: Double,
                  sameSchool      :: Bool}
```

206 Similarly, a candidate might prefer to specify the ranking of hospitals implicitly  
207 based on the livability of the city the hospital is in, reputation of the programs  
208 and their personal desire to attend a particular program. Again, these criteria  
209 are assigned appropriate weights. The applicants' model of hospital preferences  
210 is captured by the data type `HInfo`, defined as follows.

```
data HInfo = Hptl {hospitalRank  :: Rank,
                  cityLivability :: Int,
                  desirabilityScore :: Double}
```

211 Next we need to express the information in a form that supports the computation  
212 of preference lists.

### 213 3.1 Normalization and Weighting of Criteria

214 To generate rankings we normalize values of a representation type to numbers  
215 with the help of a type class `Norm`. Some examples of normalizations are shown  
216 in Figure 2b. When a type `a` is an instance of this class, the member function  
217 `norm` computes a number between 0 and 1 for any of its values. In addition to  
218 the value to be normalized, `norm` takes an optional normalization constant which  
219 acts as a bound on the range of values.

220 For example, a score of 80 in an exam out of 100 can be normalized as  $\frac{80}{100} =$   
221 0.8. We would represent this normalization as `80 `outOf` 100` (which is syntactic  
222 sugar for `norm (80, Just 100)`). As a different example, the type `Rank` was used  
223 earlier for representing relative preferences. Its `Norm` instance simply conveys that  
224 a numerically lower rank corresponds a higher preference and vice versa. Thus,  
225 we can simply use reciprocal values without the needing a normalizing constant.

226 The type class also provides a `components` function, which provides a list of  
227 normalized values corresponding to the various arguments of a constructor of an  
228 abstract data type. As can be seen in the class definition, once we have defined  
229 `components` for a data type, the normalized values can easily be inferred from  
230 it. The function `outOf` and `only` are instances of the `norm` function specialized  
231 for the cases when normalizing constants are required and when they are not,  
232 respectively.

233 With the help of `Norm` we can define the normalization for the applicant and  
234 hospital preference representations as follows.

```
instance Norm AInfo where
  components (Appl e x i c) = [e `outOf` 800, x `outOf` 10, i `outOf` 10, only c]

instance Norm HInfo where
  components (Hpt1 h c d) = [only h, c `outOf` 10, d `outOf` 5]
```

235 However, before we compute the preferences using the normalization of repre-  
236 sentation types, we need to address the situation that applicants or hospitals  
237 may weight criteria differently.

238 To that end MATCHMAKER provides a class `Weights` shown in Figure 2b,  
239 which can be used to assign different weight profiles for various criteria corre-  
240 sponding to different constructors of a type `a`. Then we amend the definition  
241 of the `Preference` type class (shown in Figure 2b) by adding a class constraint  
242 which specifies that its first type argument should also be a member of the `Weight`  
243 class, which allows us to generate rankings for various hospitals and applicants  
244 using different distributions of the criteria weights.

245 The weight distributions of the criteria for various hospitals and applicants  
246 are specified as instances of the `Weight` class. We see that Mercy assigns greater  
247 importance to exam and interview scores than to previous work experiences  
248 compared to the other hospitals. Moreover, in contrast to other hospitals, Mercy  
249 gives some weight to whether or not applicants have studied there previously.

```

instance Weights Hospital where
  weights Mercy = [0.3,0.3,0.3,0.1]
  weights _     = [0.2,0.2,0.6,0.0]

```

250 For applicants we assume that they all use the same weights for the various  
 251 criteria.

```

instance Weights Applicant where
  weights = forall [0.2,0.2,0.6]

```

## 252 3.2 Representational Rankings in Action

253 Now we can derive rank from preference representations. Specifically, we can  
 254 replace the third argument of the `Preference` type class, `Rank`, with `AInfo` and  
 255 `HInfo`, allowing us to record the preferences for hospitals and applicants, re-  
 256 spectively. Before we look at the actual preference encodings of applicants, we  
 257 observe that values of some criteria remain unchanged for the different appli-  
 258 cants. For example, rankings of the hospitals and the livability of the cities they  
 259 are located in are not applicant dependent but intrinsic to the hospitals and  
 260 cities themselves. We can exploit this fact to factor out this shared information,  
 261 which can then be used by all applicants. The function `hProfile` constructs a  
 262 hospital/city profile for each hospital as a partial `HInfo` value with fixed ranking  
 263 and livability score information but still unassigned desirability scores `DScore`  
 264 (which is a type synonym for `Double`).

```

hProfile :: Hospital -> DScore -> HInfo
hProfile Mercy   = Hptl (Rank 2) 9
hProfile City    = Hptl (Rank 1) 10
hProfile General = Hptl (Rank 3) 8

```

265 Next we represent the desirability scores of hospitals for the different applicants  
 266 in the form of an `Info` value. Of course, it may be the case that applicants may  
 267 use different sources for getting the ranking and livability information resulting  
 268 in non-uniform rankings of hospitals and livability scores of cities. In such a case,  
 269 we could have two additional `Info` values, one each for rank and livability, similar  
 270 to what we have for the desirability scores. However, for our current example,  
 271 we consider them to be uniform.

```

desirability :: Info Applicant Hospital DScore
desirability =
  info [Arthur --> [City --> 3],
        Sunny  --> [Mercy --> 4,City --> 3],
        Joseph --> [Mercy --> 1,City --> 5,General --> 4],
        Latha  --> [Mercy --> 5,City --> 1,General --> 1],
        Darrius--> [Mercy --> 5,City --> 5,General --> 4]]

```

272 We can combine the fixed and variable criteria values to generate the overall  
 273 representation of applicants' preferences using the `completedWith` combinator.  
 274 As the type of `completedWith` (shown in Figure 3) indicates, it takes a function  
 275 with output type `d` and an `Info` value with type `c` as its third type argument  
 276 representing the value type of the variable criterion. It returns as output a com-  
 277 pleted `Info` value for matching set `a` with respect to `b` using the type `d`.

```

instance Preference Applicant Hospital HInfo where
  gather = hProfile `completedWith` desirability

```

```

zipInfo  :: (Ord2 a b) => Info a b c -> Info a b d -> Info a b (c,d)
zipInfo2 :: (Ord2 a b) => Info a b c -> Info a b d -> Info a b e -> Info a b (c,d,e)

completedWith  :: Ord a => (b -> c -> d) -> Info a b c -> Info a b d
completedWith2 :: Ord a => (b -> c -> d -> e) -> Info a b (c,d) -> Info a b e

```

Fig. 3: Combinators for combining `Info` values and generating them from profiles.

278 We can represent the preferences for hospitals in a similar way. Again, we begin  
 279 by defining the profile of applicants `aProfile` for the fixed information, which  
 280 includes the applicants' exam scores and their work experience.

```

aProfile :: Applicant -> IScore -> SStatus -> AInfo
aProfile a = case a of
    Arthur -> Appl 700 2
    Sunny -> Appl 720 2
    Joseph -> Appl 750 1
    Latha -> Appl 650 5
    Darrius-> Appl 790 2

```

281 This leaves applicants' hospital-dependent attributes, such as interview scores  
 282 `IScore` and prior student status `SStatus` at a hospital, to be filled in by the  
 283 individual hospitals. The interview scores of applicants at various hospitals are  
 284 recorded again in a corresponding `Info` value.

```

interview :: Info Hospital Applicant IScore
interview = info
    [Mercy --> [Joseph --> 8,Darrius --> 9],
     City --> [Arthur -->10,Sunny --> 9,Joseph --> 4,Latha --> 6,Darrius--> 10],
     General --> [Arthur --> 9,Joseph --> 8,Latha --> 5,Darrius --> 10]]

```

285 Similarly, the student status of applicants at a given hospital is also represented  
 286 by an `Info` value.

```

school :: Info Hospital Applicant SStatus
school = info
    [Mercy -->[Joseph --> False,Darrius --> True],
     City -->[Arthur --> True,Sunny --> False,Joseph --> False,Latha --> False,Darrius --> False],
     General-->[Arthur --> False,Joseph --> True,Latha --> False,Darrius --> False]]

```

287 Finally, we can combine the applicants' profiles with their interview scores and  
 288 student status information to generate an `Info` value with complete information  
 289 about students. We do this by first “zipping” together `interview` and `school`  
 290 using the `zipInfo` function, which results in an `Info` value where the interview  
 291 score and school status information for every candidate is paired up. The function  
 292 `zipInfo` is analogous to Haskell's `zip` function in that it has the effect of pairing  
 293 `Info` values. We also provide functions `zipInfo2`, `zipInfo3` and so on, for combin-  
 294 ing multiple `Info` values, corresponding to Haskell's `zip2` and `zip3`. The function  
 295 `completedWith2` is a function which takes as input a profile with two unassigned  
 296 fields and an `Info` value that contains these variable values and produces a com-  
 297 pleted `Info` value. We provide different variants of the `completedWith` function  
 298 to join multiple `Info` values.

```

instance Preference Hospital Applicant AInfo where
    gather = aProfile `completedWith2` (interview `zipInfo` school)

```

299 This completes the specification of applicant and hospital preferences. It is in-  
300 structive to see that we can get concrete rankings from our preference represen-  
301 tations. We can do so using the `ranks` function (defined in Figure 2a) as shown  
302 below. Note that the preference lists of hospitals are unchanged from Figure 1.  
303 Similarly, we can verify that the preference lists for applicants have not changed  
304 either.

```
> ranks (gather :: Info Hospital Applicant AInfo)
{City   --> [Darrius,Arthur,Sunny,Latha,Joseph] : 2,
 Mercy  --> [Darrius,Joseph] : 2,
 General --> [Darrius,Arthur,Joseph,Latha] : 2}
```

305 The stable matchings can be generated in the same way as we did with explicit  
306 rankings.

```
> twoWay :: Match Hospital Applicant
{City   --> [Arthur,Darrius],
 Mercy  --> [],
 General --> [Latha,Joseph]}
```

307 Since the inferred preference lists for applicants and hospitals didn't change, the  
308 stable matchings don't change either.

## 309 4 Evolution and Analysis of Matches

310 So far we have seen matching problems with a fixed initial set of agents. Lets  
311 assume that, in the context of our example, some hospitals or applicants de-  
312 cide to amend their preferences or that maybe some hospitals or applicants are  
313 added late in the NRMP cycle and need to be accommodated in the match. The  
314 straightforward thing to do would be to manually modify the preference lists and  
315 rerun the matching algorithm on this amended list. Not only is this approach  
316 prone to errors during the update, we would also lose track of the history of the  
317 different stages of the process, which might be of interest to see how changes  
318 in the data lead to changes in matches. An alternative is to keep the original  
319 and amend it using functions provided by the DSL. This approach makes the  
320 changes explicit, allowing users to track the evolution of data and corresponding  
321 matchings. The type signatures for some of the relevant functions for these tasks  
322 are shown in Figure 4.

### 323 4.1 Updating Ranks and Adding Agents

324 Assume that a new applicant Bob is added to the matching process. Like other  
325 applicants, Bob will have his preference list of hospitals. Hospitals will also need  
326 to accommodate him in their preference lists. Let's further assume that City  
327 decides not to rank him. Situations like this are of special interests to game  
328 theorists who are interested in finding out how the addition of a new applicant  
329 or a hospital might change the resulting match. For example, is it more favorable  
330 to the applicants or the hospitals? In this section we look at how `MATCHMAKER`  
331 can be used to support such investigations.

332 We begin by updating the `Applicant` data type to include the `Bob` constructor.

333

```

modWithRanks :: Ord2 a b => Info a b Rank -> (a,[b]) -> Info a b Rank
modWithInfo  :: Ord2 a b => Info a b c -> Info a b c -> Info a b c
modWithRow   :: Ord2 a b => Info a b c -> (a,[(b,c)]) -> Info a b c

updateWithRow :: Ord2 a b => Info a b c -> (a,[(b,c)]) -> Info a b c
updateWithInfo :: Ord2 a b => Info a b c -> Info a b c -> Info a b c
updateWithInfos :: Ord2 a b => Info a b c -> [Info a b c] -> Info a b c

modWithRanksDef :: (Ord2 a b,Preference a b Rank) => (a,[b]) -> Info a b Rank
...

data CompMatch a b = CompMatch {unCompMatch :: [(a,[b],[b])]}

diffMatch :: Eq2 a b => Match a b -> Match a b -> CompMatch a b
twoWayDiff :: Info a b c -> Info a b c -> CompMatch a b

```

Fig. 4: Combinators to modify encodings and compare results.

```

data Applicant = Arthur | Sunny | Joseph | Latha | Darrius | Bob

```

334 We can update the preference list of applicants by adding Bob’s preferences  
335 using the `modWithRanks` function. It takes as input the original preference list of  
336 applicants as well the new applicant to be added with his preference list. The  
337 function `gather` provides the original encoding of the preferences for applicants.

```

updatedAppl = gather `modWithRanks` (Bob --> [Mercy, City, General])

```

338 We also update the preference lists for hospitals. Note how we can chain to-  
339 gether multiple updates. A difference between the two values `updatedAppl` and  
340 `updatedHosp` is that while the former creates a new record for Bob, the latter  
341 simply updates the already existing preference lists for Mercy and General.

```

updatedHosp = gather `modWithRanks` (Mercy --> [Darrius, Bob, Joseph])
              `modWithRanks` (General --> [Bob, Darrius, Arthur, Joseph, Latha])

```

342 When we need to modify the preference lists of multiple agents, rather than  
343 making one change at a time by chaining together multiple `modWithRanks` calls,  
344 it is more convenient to collect all the changes in an `Info` value and update  
345 the original encoding with it in one go. This can be done with the `modWithInfo`  
346 function, as shown below. The updated preferences of Mercy and General are  
347 stored in an `Info` value called `deltaInfo`, which can then be used to update  
348 the original preference encoding of the applicants. Note that since City doesn’t  
349 appear in `deltaInfo`, its preferences are not changed in `updatedHosp`.

```

deltaInfo = choices [Mercy --> [Darrius, Bob, Joseph],
                    General --> [Bob, Darrius, Arthur, Joseph, Latha]]

updatedHosp = gather `modWithInfo` deltaInfo

```

350 The function `modWithInfo` is useful for various reasons. When the number of  
351 elements being matched is large, we can keep the original data and the intended  
352 changes separate. If we need to make iterative changes, this approach keeps track  
353 of the changes performed in each iteration. We can also contemplate alternative  
354 changes to the data. We also have a `modWithInfos` combinator, which can be  
355 used to modify the original data with a list of iterative changes stored as `Info`  
356 values themselves. For example, the following expression modifies the data by  
357 four updates `i1`, `...`, `i4`.

```
updated = gather `modWithInfos` [i1,i2,i3,i4]
```

358 If at any point we need to undo some of the changes, we can simply remove the  
359 corresponding `Info` value from the list.

360 Now that we have the amended preference lists for hospitals and applicants,  
361 we can use them to get new matchings using the `twoWayWithPref` function, which  
362 was introduced in Section 2.2.

```
> twoWayWithPref updatedHosp updatedAppl  
{City --> [Arthur,Darrius],  
General --> [Latha,Joseph],  
Mercy --> [Bob]}
```

363 Notice how the matching is different from the original matchings, repeated here  
364 for convenience.

```
> twoWay :: Match Hospital Applicant  
{City --> [Arthur,Darrius],  
General --> [Latha,Joseph],  
Mercy --> []}
```

365 Clearly, Mercy has benefited by gaining a resident. While figuring out the differ-  
366 ence was trivial in our current example, spotting changes in even a moderately  
367 large example is more difficult. To do so systematically, we provide a function  
368 called `diffMatch`, which compares two `Match` values and reports the difference  
369 between the two matching. In our current example we obtain the following.

```
> diffMatch twoWay (twoWayWithPref updatedHosp updatedAppl)  
{Mercy --> [] => [Bob]}
```

370 The result `Mercy --> [] => [Bob]` shows that that Mercy went from not having  
371 any resident in the original match to having Bob in the updated match. An  
372 interesting thing to note here is that even though we didn't annotate the type  
373 of the first argument `twoWay`, it can be inferred from the type of the second  
374 argument of `diffMatch`.

375 What can we say about the performance of various hospitals and applicants  
376 in the updated match, compared to the original match? Intuitively, it seems that  
377 most hospitals, namely City and General, have performed as well as they did  
378 before, while Mercy has improved its performance. Similarly, it seems that no  
379 applicants have performed worse than the original match. Do these observations  
380 always hold? Game theory tells us that no hospital will be worse off and some  
381 hospitals are better off compared the original match [16, Theorem 2.26]). At the  
382 same time, none of the original applicants are better off, while some can be worse  
383 off than in the original match. In any case, MATCHMAKER can be employed as  
384 tool for gaining a deeper understanding of a wide range of matching scenarios.

## 385 4.2 Updating Representational Ranks

386 Assume that we wanted to update the representational ranks of our example  
387 from Section 3.2. More concretely, suppose Mercy wanted to add Sunny and  
388 Arthur and City wanted to add Sunny to its preference list. They only need to  
389 provide the interview scores and school status for the applicants as the other

390 information can be obtained from the applicants' profiles. The interview scores  
 391 can be updated for the two hospitals using the `updateWithRow` combinator, which  
 392 takes an `Info` value to be updated along with the information to update it  
 393 with. An entry such as `City --> [Sunny --> 9]` indicates that City assigns a  
 394 interview score of 9 to Sunny, which is then appended to its already existing score  
 395 assignments for other applicants. The function `updateWithRow` can be chained  
 396 together to update the records for multiple hospitals.

```
interview1 = interview `updateWithRow` (City --> [Sunny --> 9])
              `updateWithRow` (Mercy --> [Sunny --> 8, Arthur --> 8])
```

397 And the school status also needs to be updated.

```
school1 = school `updateWithRow` (Mercy --> [Sunny --> True, Arthur --> False])
              `updateWithRow` (City --> [Sunny --> False])
```

398 Again, we also have the option to collect all changes in an `Info` value, which is  
 399 then used by the `updateWithInfo` combinator.

```
deltaInterview = info [Mercy --> [Sunny --> 8, Arthur --> 8], City --> [Sunny --> 9]]
interview1 = interview `updateWithInfo` deltaInterview
```

400 Finally, we can use the modified interview scores and school status information  
 401 to update the preferences for hospitals.

```
updatedHosp = aProfile `completedWith2` (interview1 `zipInfo` school1)
```

402 The changed data leads to the following preference lists for various hospitals.

```
> ranks updatedHosp
{Mercy --> [Darrius, Sunny, Arthur, Joseph] : 2,
 City --> [Darrius, Arthur, Sunny, Latha, Joseph] : 2,
 General --> [Darrius, Arthur, Joseph, Latha] : 2}
```

403 We can now generate the updated match using the `twoWayWithPref` function.  
 404 But perhaps it will be more interesting to see how this matching differs from  
 405 the original match. As shown, the only difference in the two matchings is that  
 406 Mercy which was not assigned a resident initially, now has Sunny assigned to it.  
 407

```
> twoWayDiff updatedHosp gather
{Mercy --> [] => [Sunny]}
```

## 408 5 Other Matching Problems

409 In addition to the two-way stable matching problem, `MATCHMAKER` also allows  
 410 for the modeling of other interesting matching problems like one-sided matchings,  
 411 one-sided matching with exchange, and same-set matchings, which we briefly  
 412 discuss in this section. The various types and function definitions used in this  
 413 section are shown in figure 6.

$P_1$	$P_2$	$P_3$	$P_4$
Bob	Alice	Alice	Alice
Dan	Dan	Bob	Bob
Dillon	Dillon	Dillon	Dan

(a) Preference lists Donors to Patients

```

data Donor = Alice | Bob | Dan | Dillon
data Patient = P1 | P2 | P3 | P4

instance Preference Patient Donor Rank where
gather = choices [P1 --> [Bob,Dan,Dillon],
                  P2 --> [Alice,Dan,Dillon],
                  P3 --> [Alice,Bob,Dillon],
                  P4 --> [Alice,Bob,Dan]]

```

(b) Encoding the example in DSL.

Fig. 5: Assigning donors to patients: Bipartite matching with one-sided preferences.

## 414 5.1 Bipartite Matching With One-Sided Preferences

415 The first important example of a one-sided matching problem is known as the  
416 *house allocation problem* in the economics literature. In this type of matching  
417 only the elements in the source set have a preferences for the elements in the  
418 target set. The preferences of the target sets are not taken into account. Some  
419 of its applications have been allocating graduates to trainee positions, students  
420 to projects, professors to offices, and clients to servers.

421 As a concrete example, let us consider the the problem of selecting kidney  
422 donors for various transplant patients. Assume that the donors are altruistic and  
423 don't care who the kidney goes to. Patients on the other hand have a preference  
424 over the kidneys: a good kidney for a patient depends on the tissue compatibility  
425 of the donor-recipient pair as well as the donor's age and their overall health  
426 condition. Thus, the transplant team of a patient may have a ranked preference  
427 list of donors. Figure 5a shows patients with their preference lists.

428 Formally, the donor assignment problem is a three tuple  $(T, D, P)$  where  
429  $T = \{t_1, t_2, \dots, t_k\}$  is a finite set of transplant patients and  $D = \{d_1, d_2, \dots, d_n\}$   
430 is a finite set of donors.  $P$  is a preference map such that preference of each  
431 patient  $t \in T$  is represented by a ordered list of preferences  $P(t)$  on set  $D$ . We  
432 assume that each patient has a quota of 1, that is, they can be assigned just  
433 one donor. A matching  $\mu : T \rightarrow D$  in this case is a partial function that assigns  
434 every patient to 1 donor.

435 We can represent the patient-donor example with the machinery already de-  
436 veloped for two-sided matching. Figure 5 shows an encoding of the problem using  
437 explicit ranks. In a more realistic setting, the agency tasked with performing the  
438 match might prefer to rank the donors using meaningful representation such as  
439 age and the blood and tissue compatibility between the donor-patient pair.

440 How do we assign donors to the patients based on their preferences? The  
441 strategy we use here is the so-called *serial dictatorship mechanism* [1]. It is a  
442 straightforward greedy algorithm that takes each patient in turn and assigns  
443 them to the most preferred available donor on their preference list. The order in  
444 which the patients are processed will, in general, affect the outcome. In appli-

445 cations where elements have a quota of  $n$ , they are assigned to  $n$  objects when  
 446 their turn comes for processing. For our example here, we expect that a match-  
 447 ing agency will come up with an order of processing based on factors such as  
 448 the urgency of a patient’s situation, their age, or their time on the waiting list.  
 449 The function `oneWayWithOrder` performs serial dictatorship with a given order as  
 450 shown below where patient  $P_3$  gets its first choice donor,  $P_4$  gets its first choice  
 451 amongst the remaining donors, and so on.

```
> oneWayWithOrder [P3,P4,P2,P1] :: Match Patient Donor
{P1 --> [Dillon],P2 --> [Dan],P3 --> [Alice],P4 --> [Bob]}
```

452 Oftentimes users might prefer that the matching function infer a preferred order  
 453 based on position of the constructor in the data definition for donors, that is,  
 454 the `Donor` data definition implies an order of `[P1,P2,P3,P4]`. The function `oneWay`  
 455 generates a one-way match with this implicit order.

```
> oneWay :: Match Patient Donor
{P1 --> [Bob],P2 --> [Alice],P3 --> [Dillon],P4 --> [Dan]}
```

456 Finally, there is also a third variant of the function `oneWayWithPref` that takes  
 457 explicit preference encoding like its counterpart `twoWayWithPref`.

458 As we can see, these two matches are different because they are generated  
 459 using different orders. Is one better than the other? What is the best possi-  
 460 ble match among the various possibilities? Manually comparing one match with  
 461 another is cumbersome because for every patient we have to look at the two  
 462 matchings and compare the relative ranks of the two donors in that patient’s  
 463 preference list. This task is simplified by the function `diffRanks`, which com-  
 464 pares the ranks of the two matchings using a type called `CompRanks`. This type  
 465 represents for every agent the element assigned to them in those matchings  
 466 as well the elements’ ranks for comparison. In the following expression we use  
 467 `x = oneWayWithOrder [P3,P4,P2,P1]`.

```
> diffRanks oneWay x :: CompRanks Patient Donor
{P1 --> Bob : 1 > Dillon : 3, P2 --> Alice : 1 > Dan : 2,
 P3 --> Dillon : 3 < Alice : 1, P4 --> Dan : 3 < Bob : 2}
```

468 It turns out that the first match is advantageous for patients  $P_1$  and  $P_2$ , whereas  
 469 the second match is advantageous for patients  $P_3$  and  $P_4$ . Informally, a matching  
 470 is *Pareto optimal* if there is no other matching in which some patient is better  
 471 off, whilst no patient is worse off. It is used as a metric to compare the quality of  
 472 outcomes in game theoretic matchings. It turns out that the deceptively simple-  
 473 looking *serial dictatorship* algorithm results in Pareto optimal matchings, which  
 474 implies that for any two matchings there are some patients for whom one match  
 475 is better and for some the second match is better, that is, there doesn’t exist an  
 476 unique best match.

## 477 5.2 Bipartite Matching With One-Sided Preferences and Exchange

478 We assumed the presence of altruistic donors in our last example. However, kid-  
 479 neys are valuable commodities, and altruistic donors alone can’t fulfill the vast  
 480 demand for it. A more realistic scenario is a family member or a friend donating

```

class Preference a b c => Exchange a b where
  endowment :: Match a b

type SameSetMatch a = Maybe (Match a a)

data CompRanks a b = CompRanks {
  unCompRanks :: [(a, [(b, Rank)], [(b, Rank)])]}

oneWay :: (Preference a b c, Set2 a b, Norm c) => Match a b

oneWayWithOrder :: (Preference a b c, Set2 a b, Norm c) =>
  [a] -> Match a b
oneWayWithPref :: (Preference a b c, Set2 a b, Norm c) =>
  Info a b c -> Match a b

trade :: (Preference a b c, Set2 a b, Norm c) => Match a b
sameSet :: (Preference a a b, Set a, Norm b) => SameSetMatch a

spas :: (Preference a b c, Preference d a e,
  Preference d b Rank) => Match a b

diffRanks :: (Eq2 a b, Preference a b c, Set2 a b, Norm c) =>
  Match a b -> Match a b -> CompRanks a b

```

Fig. 6: Some type and function definitions for various matching problems

481 one of their kidneys to a loved one. However, sometimes this donation may not  
482 happen due to reasons like tissue or blood group incompatibility. An elegant so-  
483 lution was developed in the field of economics. Suppose  $(d_1, r_1)$  and  $(d_2, r_2)$  are  
484 two donor-receiver pairs such that  $d_i$  wants to donate to  $r_i$  but can't do so. How-  
485 ever, if  $d_1$  could donate to  $r_2$  and  $d_2$  to  $r_1$ , then both the patients would be able  
486 to receive kidneys. This could be easily scaled to multiple pairs generating large  
487 numbers of compatibility pairs. The actual *kidney exchange mechanism* [17] is a  
488 little more complicated, but the exchange between multiple donor-receiver pairs  
489 is at the heart of it. This exchange characterizes our next matching algorithm,  
490 the so-called *top trading cycle* (TTC) matching mechanism for one-way match-  
491 ing where every element has an initial endowment and a preference list [18]. The  
492 resulting match takes both of these into account.

493 We consider the same patient-donor example we considered in the last sec-  
494 tion. At the start, some donor, presumably family or friends willing to donate  
495 a kidney, is assigned to each patient. These initial set of donors are sometimes  
496 also called the *initial endowment*, or just *endowment* of a patient. Assume that  
497 patients  $P_1, \dots, P_4$  are endowed with Bob, Dan, Alice, and Dillon, respectively,  
498 such that all the patients are compatible with the donors they are endowed  
499 with. In this case, TTC tries to find out if the patients can do better than the  
500 donor they are assigned to, based on their preference lists. We start by represent-  
501 ing endowments for which we define the multi-parameter type class **Exchange**,  
502 which has a **Preference** class constraint (see Figure 6). The instance definition  
503 of **Exchange** for our example is as follows.

```

instance Exchange Patient Donor where
  endowment = assign [P1 --> Bob,
                     P2 --> Dillon,
                     P3 --> Alice,
                     P4 --> Dan]

```

```

data Student = Charlie | Peter | Kelly | Sam

instance Preference Student Student Rank where
  gather = choices [Charlie--> [Peter,Sam,Kelly],
                   Peter  --> [Kelly,Sam,Charlie],
                   Kelly  --> [Peter,Charlie,Sam],
                   Sam    --> [Charlie,Kelly,Peter]]

```

Fig. 7: A stable roommate example in MATCHMAKER.

504 Now we can use the function `trade` provided by the DSL to generate the match-  
 505 ing.

```

> trade :: Match Patient Donor
{P1 --> [Bob], P2 --> [Dan], P3 --> [Alice],
 P4 --> [Dillon]}

```

506 Did any patient gain as a result of the change? We can use the `diffRanks` function  
 507 we saw in the previous section to find out. We discover that patients  $P_2$  and  $P_4$   
 508 do indeed profit by exchanging their donors.

```

> diffRanks endowment trade :: CompRanks Patient Donor
{P2 --> Dillon : 3 < Dan : 2, P4 --> Dan : 3 < Dillon : 2}

```

### 509 5.3 Same-Set Matching

510 This variation of the problem is the so-called *Stable roommate problem* [6, 9]  
 511 where the source and the target sets being matched are the same. For example,  
 512 a set of students living in the dormitory can supply a ranked preference list of  
 513 other students they want to be roommates with. An example is shown in Figure  
 514 5b. We can obtain a stable matching of roommates using Irving’s algorithm [8].  
 515 In order to capture the fact that source and target sets are the same, we define a  
 516 type synonym `SameSetMatch` that assigns the same type `a` to both the source and  
 517 the target sets in the `Match` type (see Figure 6). Even though same-set matchings  
 518 are stable matching problems like the bipartite two-sided matching problems,  
 519 they are different in that a stable match always exists for the former, whereas it  
 520 may not always exist for the latter. This fact is reflected by the `Maybe` constructor  
 521 in the type definition of `SameSetMatch`. Finally, we can generate the the same-  
 522 set matching using the `sameSet` function, which produces for our example from  
 523 Figure 7 the following result.

```

> sameSet :: SameSetMatch Student
Just {Charlie --> [Sam], Peter --> [Kelly]}

```

## 524 6 Related Work

525 *Matching* [20] is a library for Python that allows users to encode simple matching  
 526 problems in a straightforward manner. An issue with the library is that all  
 527 the encoding are done using strings, which makes error handling difficult and

528 thus complicates the maintenance and debugging larger examples. In comparison  
529 MATCHMAKER avails the strongly typed feature of the host language Haskell to  
530 detect the various errors in encoding.

531 Similarly, *matchingMarkets* [10] is a matching library for R. The advantage  
532 of the library is that it implements a wide variety of matching algorithms de-  
533 veloped in the matching theory. Additionally, it implements statistical tools to  
534 correct for the sample selection bias from observed outcomes in matching mar-  
535 kets, which is something that MATCHMAKER doesn't do. The library encodes  
536 the preference relation between the sets of elements being matched in the form of  
537 a matrix. While an efficient way to encode the preferences, the matrix encoding  
538 is clunky and is thus difficult to understand, update and maintain. *matchingR*  
539 [19] is another stable matching library for R and C++, which uses matrices  
540 to encode the preference relations and thus suffers from the same problems as  
541 *matchingMarkets*.

542 MATCHMAKER allows users to specify their preferences more abstractly in  
543 terms of attributes that they understand, while all the previous libraries only  
544 allow specification of preference in terms of ranks. Additionally, none of these  
545 libraries offers either the primitives for systematic modification of representations  
546 or primitives to compare and contrast different matchings.

547 Matching problems can also be solved using constraint programming [13, 5] or  
548 SMT solving [3]. Moreover, integer linear programming can be used to solve NP-  
549 hard stable marriage problems, including ones with ties and incomplete lists as  
550 well as the many-to-one generalization [2]. While powerful, a potential downside  
551 is that encoding matching problems as constraints might be challenging for users.  
552 In contrast, MATCHMAKER facilitates high-level representations of matching  
553 problems and can thus be used without any specialized knowledge.

## 554 7 Conclusions

555 MATCHMAKER is an embedded DSL in Haskell for expressing, solving, and an-  
556 alyzing game-theoretic matching problems. Our implementation leverages ad-  
557 vanced type system features of Haskell to facilitate high-level representations of  
558 matching problems, expressed in terms of domain elements. MATCHMAKER also  
559 supports the maintenance and evolution of the problem representation and pro-  
560 vides some limited support for analyzing computed results, making it a useful  
561 tool for end users as well as game theorists.

562 What can we learn from our DSL design? The most important take-away  
563 message for us is the impact of the strong typing approach. In particular, bas-  
564 ing matching on data constructors instead of strings allows the detection of  
565 many potential errors during compile time. Another benefit is obtained from the  
566 use of multi-parameter type classes. For example, a definition such as `instance`  
567 `Preference Applicant Hospital Rank` serves as a reminder that a user is stating  
568 preferences of applicants for hospitals using the Rank type, supporting a mental  
569 model of the problem, which has a similar guiding effect as function signatures  
570 have for the implementation of complex functions.

571 **References**

- 572 1. Bogomolnaia, A., Moulin, H.: A new solution to the random assignment problem.  
573 *Journal of Economic Theory* **100**(2), 295–328 (2001)
- 574 2. Delorme, M., García, S., Gondzio, J., Kalcsics, J., Manlove, D., Pettersson, D.:  
575 Mathematical models for stable matching problems with ties and incomplete lists.  
576 *European Journal of Operational Research* **277**(2), 426–441 (2019)
- 577 3. Drummond, J., Perrault, A., Bacchus, F.: Sat is an effective and complete method  
578 for solving stable matching problems with couples. In: *Proceedings of the 24th*  
579 *International Conference on Artificial Intelligence*. p. 518–525. IJCAI’15, AAAI  
580 Press (2015)
- 581 4. Gale, D., Shapley, L.S.: College admissions and the stability of marriage. *The*  
582 *American Mathematical Monthly* **69**(1), 9–15 (1962)
- 583 5. Gent, I.P., Irving, R.W., Manlove, D., Prosser, P., Smith, B.M.: A constraint pro-  
584 gramming approach to the stable marriage problem. In: *Proc. of the 7th Inter-*  
585 *national Conference on Principles and Practice of Constraint Programming*. p.  
586 225–239. CP ’01, Springer-Verlag (2001)
- 587 6. Gusfield, D.: The structure of the stable roommate problem: Efficient represen-  
588 tation and enumeration of all stable assignments. *SIAM Journal on Computing*  
589 **17**(4), 742–769 (1988)
- 590 7. Gusfield, D., Irving, R.W.: *The Stable Marriage Problem: Structure and Algo-*  
591 *rithms*. MIT Press, Cambridge, MA, USA (1989)
- 592 8. Irving, R.W.: An efficient algorithm for the “stable roommates” problem. *Journal*  
593 *of Algorithms* **6**(4), 577–595 (1985)
- 594 9. Irving, R.W., Leather, P.: The complexity of counting stable marriages. *SIAM*  
595 *Journal on Computing* **15**(3), 655–667 (1986)
- 596 10. Klein, T., Aue, R., Giegerich, S., Sauer, A.: *matchingMarkets: Analysis of Stable*  
597 *Matchings in R* (2020), <https://matchingmarkets.org/>
- 598 11. Manlove, D.F.: *Algorithmics of Matching Under Preferences*. World Scientific  
599 (2013)
- 600 12. NRMP: National Resident Matching Program (2022),  
601 <https://www.nrmp.org/intro-to-the-match/how-matching-algorithm-works/>
- 602 13. Prosser, P.: Stable roommates and constraint programming. In: *CPAIOR* (2014)
- 603 14. Roth, A.E.: Deferred acceptance algorithms: History, theory, practice, and open  
604 questions. Working Paper 13225, National Bureau of Economic Research (2007)
- 605 15. Roth, A.E., Peranson, E.: The redesign of the matching market for american physi-  
606 cians: Some engineering aspects of economic design. *American Economic Review*  
607 **89**(4), 748–780 (September 1999)
- 608 16. Roth, A.E., Sotomayor, M.A.O.: *Two-Sided Matching: A Study in Game-Theoretic*  
609 *Modeling and Analysis*. Econometric Society Monographs, Cambridge University  
610 Press (1990)
- 611 17. Roth, A.E., Sönmez, T., Ünver, M.U.: Kidney exchange. *The Quarterly Journal of*  
612 *Economics* **119**(2), 457–488 (2004)
- 613 18. Shapley, L., Scarf, H.: On cores and indivisibility. *Journal of Mathematical Eco-*  
614 *nomics* **1**(1), 23–37 (1974)
- 615 19. Tilly, J., Janetos, N.: *matchingR: Matching Algorithms in R and C++* (2020),  
616 <https://github.com/jtilly/matchingR/>
- 617 20. Wilde, H., Knight, V., Gillard, J.: *Matching: A python library for solving matching*  
618 *games*. *Journal of Open Source Software* **5**(48), 2169 (2020)
- 619 21. Williams, K.J., Werth, V.P., Wolff, J.A.: An analysis of the resident match. *New*  
620 *England Journal of Medicine* **304**(19), 1165–1166 (1981)