A Comparison of Several Evolutionary Heuristics for the Frequency Assignment Problem

Carlos Cotta, José M. Troya

Dpto. de Lenguajes y Ciencias de la Computación, Univ. de Málaga
Campus de Teatinos (3.2.49), 29071 - Málaga - Spain
{ccottap, troya}@lcc.uma.es

Abstract. The Frequency Assignment Problem (FAP) is a very important problem of practical interest. This work compares several evolutionary approaches to this problem, based both in the *forma analysis* framework, and in the *decoder* paradigm. All approaches are studied from the point of view of two different quality measures of assignments: the number of distinct frequencies, and the frequency span. It is shown that using decoders as embedded heuristics is more adequate than performing a direct search in the feasible solution space. Furthermore, despite the apparent symmetry of the problem, a recombination operator based on multiple-emitter-to-frequency preservation performs better than focusing on multiple-frequency-to-emitter preservation.

1 Introduction

The term Frequency Assignment Problem (FAP) comprises a number of optimization problems of great difficulty (NP-hard in general). Although presented under different flavors, all FAPs essentially consist of finding an assignment of a set of frequencies to a set of emitters fulfilling some specific constraints (e.g., avoiding interference between closely located emitters). The actual proliferation of cellular phone networks, local television stations, etc. clearly underpins the practical interest of these problems.

The above mentioned NP-hardness of most FAPs imply that exact techniques are inherently limited for solving these problems. For this reason heuristic techniques such as tabu search, simulated annealing or genetic algorithms (GAs) are frequently used for the resolution of FAPs [3, 4, 8]. This work focuses on the application of GAs for this purpose. To be precise, we compare two evolutionary approaches to this problem, a direct search in feasible space via specifically designed operators, and an indirect search via permutation decoders.

The remainder of the paper is organized as follows. First, Section 2 provides a formal definition of the FAPs considered in this work. Next, the two approaches are presented in Section 3, describing some different variants of each one. Subsequently, empirical results are reported in Section 4. Finally, some conclusions are extracted and future work is outlined in Section 5.

2 Frequency Assignment Problems

As mentioned in the previous section, there exist a number of FAP variants, so it is necessary to give a precise definition of the particular FAPs considered in this work. As with any optimization problem, three elements must be given in order to define a FAP: a characterization of problem instances, a characterization of problem solutions, and a quality measure.

Definition 1 (FAP Instance). *An instance of the FAP is a tuple* $FAP(\mathcal{E}, \mathcal{F}, \mathcal{D}, \mathcal{R}, \mathcal{I})$ *where*

- $\mathcal{E} = \{e_1, \cdots, e_n\}$ *is a set of* emitters.
- $\mathcal{F} = \{f_1, \cdots, f_m\}$ *is a set of available* frequencies.
- $\mathcal{D} : \mathcal{E} \times \mathcal{E} \to \mathbb{R}$ *is a function such that* $D(e, e')$ *is the distance between emitters* e *and* e'.
- $\mathcal{R} : \mathcal{E} \to \mathbb{N}$ *is a function such that* $\mathcal{R}(e)$ *is the number of frequencies required by emitter* e.
- $\mathcal{I} : \mathbb{R} \to \mathbb{N}$ *is a function such that* $\mathcal{I}(d)$ *is the frequency separation needed to avoid interference between two emitters separated by a distance* d.

According to this definition, it is easy to see that there exist two central constraints in a FAP instance that must be satisfied: the number of frequencies assigned to an emitter must be equal to the number of frequencies it demands, and these frequencies must not interfere with frequencies assigned to other emitters. This is formalized below:

Definition 2 (FAP Solution). *A solution for a FAP instance* $I(\mathcal{E}, \mathcal{F}, \mathcal{D}, \mathcal{R}, \mathcal{I})$ *is a vector* $\alpha = \langle \alpha_1, \cdots, \alpha_n \rangle \in [2^{\mathcal{F}}]^n$ *such that*

- $|\alpha_e| = \mathcal{R}(e)$, *i.e., each emitter is assigned the number of different frequencies it demands.*
- $\forall e, e' \in \mathcal{E} \; \nexists f, f' \in \mathcal{F} \; : f \in \alpha_e, \; f' \in \alpha_{e'}, \; |f - f'| < \mathcal{I}(\mathcal{D}(e, e'))$, *i.e., no interfering frequencies are assigned to two emitters.*

For the purposes of this work, we will consider that $|\mathcal{F}|$ is high enough to allow the existence of FAP solutions as shown in the previous definition. An upper bound for the cardinality of $\mathcal{F}$ is thus $\mathcal{I}(0) \cdot \sum_{e \in \mathcal{E}} \mathcal{R}(e)$, assuming that $\mathcal{I}$ is monotonically decreasing, as usual.

A quality function must be defined now, in order to quantify the goodness of a particular FAP solution. In this work, we will consider two different quality measures. The first one is termed the *frequency span*, and is defined below.

Definition 3 (Frequency Span). *The frequency span* $F(\alpha)$ *of a FAP solution* α *is*

$$F(\alpha) = \max_{e, e' \in \mathcal{E}} \left[\max_{f \in \alpha_e, f' \in \alpha_{e'}} (|f - f'|) \right] \; , \tag{1}$$

i.e., the maximum separation between assigned frequencies.

Thus, the optimal solution with respect to this quality measure is the one that satisfies the problem constraints within the smallest frequency interval. This is important in situations in which the frequency spectrum is partitioned into disjoint compact sets (e.g., a set of frequencies per city or province), and we require to fit the frequency demand of a group of emitters within one of these sets. A related, but generally different measure of FAP solution is its *size*:

Definition 4 (Assignment Size). *The size S of a FAP solution α is*

$$S(\alpha) = |\cup_{e \in \mathcal{E}} \alpha_e| \ , \tag{2}$$

i.e., the number of different frequencies assigned to emitters in $\mathcal{E}$.

Hence, the above quality measure tries to promote frequency re-utilization (notice that this re-utilization does not necessarily result in lower frequency spans).

3 Two Evolutionary Approaches for the Frequency Assignment Problem

This section will describe two different approaches for tackling FAP instances. Both mechanisms are based on restricting the search to the feasible space, but differ in the way they achieve this. On one hand, FAP solutions can be directly manipulated during recombination and mutation. On the other hand, this manipulation can be done indirectly via a construction heuristic. These two approaches are discussed below.

3.1 Direct Manipulation in Feasible-Space

As mentioned above, the first approach consists of directly manipulating frequency assignments. It is thus necessary to define the information units that will be subject to this manipulation. Let us consider the set of equivalence relations $\Psi = \{\psi_{ef} \mid e \in \mathcal{E},\ f \in \mathcal{F}\}$, where $\psi_{ef}(\alpha, \alpha') = \texttt{TRUE}$ if, and only if, frequency f is assigned to emitter e both in α and α' or in neither of them. Subsequently, each equivalence relation ψ_{ef} induces two equivalence classes, respectively comprising solutions assigning f to e (ψ^1_{ef}) or not (ψ^0_{ef}). Each of these equivalence classes is termed a *basic forma* [6].

Ψ can be shown to be an independent set covering the feasible space, so it can be used to induce a representation of solutions, i.e., $\alpha = \langle \alpha_1, \cdots, \alpha_n \rangle \equiv \bigcap_{i=1}^{n} \bigcap_{f \in \alpha_i} \psi^1_{e_i f}$. On the basis of this representation, three different information units can be processed: partial emitter assignments (i.e., formae ϕ^F_e, $e \in \mathcal{E}$, $F \in 2^{\mathcal{F}}$, defined as $\phi^F_e = \bigcap_{f \in F} \psi^1_{ef}$), partial frequency assignments (i.e., formae η^E_f, $f \in \mathcal{F}$, $E \in 2^{\mathcal{E}}$ defined as $\eta^E_f = \bigcap_{e \in E} \psi^1_{ef}$), and single emitter-frequency assignments (i.e., single formae ψ^1_{ef}, $e \in \mathcal{E}$, $f \in \mathcal{F}$).

In order to manipulate these information units, it must be taken into account that none of them are orthogonal, although they are separable. This is formally established below.

Proposition 1. *Single emitter-frequency assignments are not orthogonal.*

Proof. *The proof is straightforward. Given two formae ψ^1_{ef} and $\psi^1_{e'f'}$, their intersection is empty if $|f - f'| < \mathcal{I}(\mathcal{D}(e, e'))$. This is true for many values of e, e', f, and f' unless $\mathcal{R}(d) = 0$ for all d (a trivial situation without interest from an optimization point of view).* □

Since partial emitter-assignments and partial frequency-assignments are defined as the intersection of single emitter-frequency assignments, it follows as a corollary that none of these units is orthogonal.

Proposition 2. *Single emitter-frequency assignments are separable.*

Proof. *It must be shown that given ψ^1_{ef} and $\psi^1_{e'f'}$ ($\psi^1_{ef} \cap \psi^1_{e'f'} \neq \emptyset$), no $\alpha \in \psi^1_{ef}$, $\alpha' \in \psi^1_{e'f'}$ exist such that $\psi^1_{ef} \cap \psi^1_{e'f'} \cap \Omega = \emptyset$, where Ω is the intersection of all basic formae common to α and α'. For this latter intersection to be empty, it must be that either f or f' (or both) interfere with some frequency assignments in Ω. But this is impossible because the existence of α and α' respectively implies that $\psi^1_{ef} \cap \Omega \neq \emptyset$ and $\psi^1_{e'f'} \cap \Omega \neq \emptyset$.* □

This separability result implies that the information common to two assignments α and α' can be respected and, simultaneously, compatible information can be assorted. Nevertheless, this is generally incompatible with forma transmission, i.e., it may be necessary to introduce some exogenous information -not present either in α or α'- in the assignment produced during recombination. This new information could be selected at random, or by means of a heuristic. In this work we consider the latter approach, which falls within the patching model termed *locally optimal completion* [7] (the precise heuristics used for this purpose will be described in the next subsection).

3.2 The Decoder Approach

The decoder approach is a completely different way of carrying out the search in the feasible region. In this case, assignments are not directly manipulated. On the contrary, some external structures are processed, being a so-called *decoder* used to translate these structures into feasible assignments in order to perform evaluation. A very typical situation is the use of permutation decoders, due to the fact that permutations are a well-known structure for which different reproductive operators are available (a good survey can be found in [2]). This is the approach considered in this work.

Before defining the particular decoders considered in this work, notice that FAPs are closely related to coloring problems, as pointed out in [5]. A FAP instance can be represented as a labeled graph $G(V, E)$, where $V \equiv \mathcal{E}$, and $(e, e', \delta) \in E \Leftrightarrow \mathcal{I}(\mathcal{D}(e, e')) = \delta$ ($\delta > 0$). A FAP solution would then be a multicoloring of the graph, such that each vertex is assigned as many colors as frequencies demands, and the colors assigned to adjacent vertices satisfy the separation constraint δ of the edge connecting them. An algorithm for obtaining such a coloring of the graph is shown in Fig. 1.

First-Available-Frequency Heuristic

1. Let $P = \langle e_{i_1}, e_{i_2}, \cdots, e_{i_n} \rangle$ be a permutation of the vertices in V.
2. For all $j \in \{1, \cdots, n\}$ do $\mathcal{A}_j \leftarrow \mathcal{F}$.
3. For all $j \in \{1, \cdots, n\}$ do
   (a) Let $\alpha_{e_{i_j}} \leftarrow \emptyset$.
   (b) For all $k \in \{1, \cdots, \mathcal{R}(e_{i_j})\}$ do
      i. Let $f \leftarrow \min_{f' \in \mathcal{A}_j} f'$.
      ii. Let $\alpha_{e_{i_j}} \leftarrow \alpha_{e_{i_j}} \cup \{f\}$.
      iii. For all $j' \in \{j, \cdots, n\}$, $(e_{i_j}, e_{i_{j'}}, \delta) \in E$ do
         $\mathcal{A}_{e_{j'}} \leftarrow \mathcal{A}_{e_{j'}} - \{f' \mid \delta > |f - f'|\}$.

Fig. 1. Pseudocode of the First-Available-Frequency heuristic.

Notice now that associated to the mentioned labeled graph, there exists a dual graph in which the vertices are frequencies, and edges are labeled with subsets of $\mathcal{E}$. In this dual graph, the edge (f, f', σ) means that frequencies f and f' cannot be simultaneously assigned to nodes $e, e' \in \sigma$. Hence, a FAP solution can be also obtained by multicoloring this graph. This can be done using the algorithm depicted in Fig. 2.

First-Available-Emitter Heuristic

1. Let $P = \langle e_{i_1}, e_{i_2}, \cdots, e_{i_n} \rangle$ be a permutation of the vertices in V.
2. For all $j \in \{1, \cdots, n\}$ do
   (a) Let $\alpha_{e_{i_j}} \leftarrow \emptyset$.
   (b) Let $\mathcal{A}_j \leftarrow \mathcal{F}$.
3. Let $T \leftarrow \sum_{e \in \mathcal{E}} \mathcal{R}(e)$.
4. Let $f \leftarrow \min_{f' \in \mathcal{F}} f'$.
5. While $T > 0$ do
   (a) For all $j \in \{1, \cdots, n\}$ do
      If $\left[|\alpha_{e_{i_j}}| < \mathcal{R}(e_{i_j})\right] \wedge (f \in \mathcal{A}_{e_j})$ then
      i. Let $\alpha_{e_{i_j}} \leftarrow \alpha_{e_{i_j}} \cup \{f\}$.
      ii. For all $(e_{i_j}, e_{i_{j'}}, \delta) \in E$ do
         $\mathcal{A}_{e_{j'}} \leftarrow \mathcal{A}_{e_{j'}} - \{f' \mid \delta > |f - f'|\}$.
      iii. Let $T \leftarrow T - 1$.
   (b) Let $f \leftarrow f + 1$.

Fig. 2. Pseudocode of the First-Available-Emitter heuristic.

Both algorithms can be used as a decoders in a permutation-based GA. This allows the utilization of classical recombination/mutation operators during the reproductive stage.

4 Experimental Results

The test suite used in this work is composed of eight 21-emitter FAP instances. These eight instances correspond to the combination of two emitter layouts and four frequency-demand vectors. The first layout is a random distribution of emitters within a 6×6 plane, and the second one is the well-known Philadelphia layout [1], based on a cellular-phone network. Both frequency-demand vectors and interference constraints are taken from [8].

The first experiments consist of a fitness-variance analysis. The goal of these experiments is estimating which of the two views of the problem (emitter-based or frequency-based) carries more significant fitness information. This is important from the perspective of both the direct approach and the decoder approach. The results of this analysis are shown in Fig. 3. As it can be seen, the fitness variance is lower (and hence the fitness information is more significant) when processing frequency-based units (i.e., manipulating partial frequency assignments, or using the First-Available-Emitter heuristic).

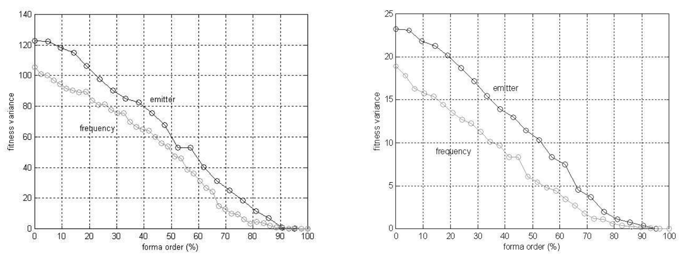

Fig. 3. Fitness variance for emitter-based units, and for frequency-based units. (Left) Frequency span (Right) Number of frequencies.

To confirm these results, the four random-layout instances are used. To be precise, the two decoder variants have been tested, using an elitist generational genetic algorithms ($popsize = 100$, $p_c = .9$, $p_m = 0.013$, $maxevals = 100.000$) utilizing ranking selection ($\eta^+ = 2.0$). Three different recombination operators have been used: cycle crossover (CX), order crossover (OX), and partially mapped crossover (PMX). In all cases, mutation is done via the *swap* operator. The results are shown in Table 1.

As it can be seen, the First-Available-Emitter heuristic is globally better that the First-Available-Frequency heuristic, confirming the hypothesis extracted from the fitness-variance analysis. Hence, the former heuristic will be used as the patching algorithm in subsequent experiments. These are done on the Philadelphia instances, using the same experimental setup mentioned above. The results are shown in Table 2.

The obtained results are conclusive. On one hand, FX (frequency crossover) performs better than EX (emitter crossover) or EFX (emitter-frequency crossover).

Table 1. Comparison of different genetic operators in the decoder approach (random-layout instances). All results correspond to series of twenty runs.

	IR1			IR2			IR3			IR4		
	CX	OX	PMX	CX	OX	PMX	CX	OX	PMX	CX	OX	PMX
	Frequency Span (mean)											
FAF	421.0	421.0	421.0	295.8	295.8	295.6	288.8	288.6	289.6	846.0	846.0	846.0
FAE	421.0	421.0	421.0	295.0	295.0	295.0	219.1	219.0	219.1	846.0	846.0	846.0
	Number of Frequencies (mean)											
FAF	301.2	299.7	300.8	261.0	258.4	259.1	220.4	220.3	220.1	594.1	595.0	596.1
FAE	295.8	296.2	298.1	258.9	260.8	263.7	220.0	220.0	220.0	590.1	594.2	593.8

This is in accordance with the previous fitness-variance analysis, and with the fact that EFX is an operator performing a strong mixture of information taken from the parents (similar to UX in binary representations). This is a detrimental property in such a constrained problem, in which assignment values closely interact. On the other hand, the decoder approach yields the overall best results. Actually, the optimal solution is found for three of the problem instances (IP1, IP3, and IP4), since the GA reaches the lower bounds given in [8]. This good performance can be partially explained by the fact that the decoder provides locally optimal solutions of high quality, difficult to achieve by random recombination. Clearly, this is a specific property of the particular decoding algorithm used in this work. In this sense, the First-Available-Emitter heuristic seems to be very appropriate to introduce problem-specific knowledge in the GA.

5 Conclusions

This work has compared two different approaches for the resolution of frequency assignment problems. The obtained results have confirmed the goodness of fitness-variance estimations in order to predict GA performance. It has been shown that, despite the apparent emitter/frequency symmetry of the problem, manipulating partial frequency assignments is more adequate that manipulating partial emitter assignments. This results also holds for two different construction heuristics used as decoders in a permutation-based GA. Furthermore, there seems to be a good interplay between the GA and the First-Available-Emitter heuristic, resulting in much better solutions than those obtained by means of *blind* recombination operators.

Future work will be directed to study other construction heuristics, as well as tackling different variants of FAPs. Over-constrained instances in which the goal is minimizing unfeasibility rather than optimizing feasibility are a specifically interesting line of future work.

Acknowledgments. This work is partially supported by CICYT under grant TIC99-0754-C03-03.

Table 2. Comparison of different genetic operators on the Philadelphia test-suite. Patching is done via the First-Available-Emitter heuristic. All results correspond to series of twenty runs.

Frequency Span												
	IP1			IP2			IP3			IP4		
Operator	best	mean	σ	best	mean	σ	best	mean	σ	best	mean	σ
EX	434	459.40	12.95	287	308.25	11.73	243	287.70	12.08	875	919.50	25.89
FX	434	451.10	8.41	269	275.65	4.40	240	244.65	1.49	871	907.55	19.55
EFX	435	459.25	13.27	296	308.00	8.14	243	270.60	8.79	885	920.50	24.07
OX	426	426.00	0.00	273	281.75	4.94	239	239.00	0.00	855	855.90	0.94
PMX	426	426.00	0.00	271	285.45	5.32	239	239.00	0.00	855	855.65	0.57
CX	426	426.25	0.43	282	288.25	3.16	239	239.00	0.00	855	857.10	1.45
Number of Frequencies												
	IP1			IP2			IP3			IP4		
Operator	best	mean	σ	best	mean	σ	best	mean	σ	best	mean	σ
EX	360	360.10	0.30	271	275.25	3.39	240	241.95	4.34	720	720.70	0.95
FX	360	360.00	0.00	270	270.00	0.00	240	240.00	0.00	720	720.00	0.00
EFX	360	363.70	2.33	273	282.95	4.66	240	249.25	9.03	720	727.40	4.86
OX	360	360.00	0.00	270	270.00	0.00	240	240.00	0.00	720	720.00	0.00
PMX	360	360.00	0.00	270	270.00	0.00	240	240.00	0.00	720	720.00	0.00
CX	360	360.00	0.00	270	270.00	0.00	240	240.00	0.00	720	720.00	0.00

References

1. L.G. Anderson. A simulation study of some dynamic channel assignment algorithms in a high capacity mobile telecommunications system. *IEEE Transactions on Communications*, COM-21:1294–1301, 1973.
2. C. Cotta and J.M. Troya. Genetic forma recombination in permutation flowshop problems. *Evolutionary Computation*, 6(1):25–44, 1998.
3. C. Crisan and H. Mühlenbein. The breeder genetic algorithm for frequency assignment. In A.E. Eiben et al., editors, *Parallel Problem Solving From Nature V - LNCS 1498*, pages 897–906. Springer-Verlag, Berlin, 1998.
4. S. Hurley, D.J. Smith, and S.U. Thiel. FASoft: a system for discrete channel frequency assignment. *Radio Science*, 32:1921–1939, 1997.
5. R.A. Murphey, P.M. Pardalos, and M.G.C. Resende. Frequency assignment problems. In D.-Z. Du and P. M. Pardalos, editors, *Handbook of combinatorial optimization*, volume 3. Kluwer Academic Publishers, 1999.
6. N.J. Radcliffe. Equivalence class analysis of genetic algorithms. *Complex Systems*, 5:183–205, 1991.
7. N.J. Radcliffe and P.D. Surry. Fitness variance of formae and performance prediction. In L.D. Whitley and M.D. Vose, editors, *Foundations of Genetic Algorithms III*, pages 51–72, San Mateo CA, 1994. Morgan Kauffman.
8. C. Valenzuela, S. Hurley, and D. Smith. A permutation based genetic algorithm for minimum span frequency assignment. In A.E. Eiben et al., editors, *Parallel Problem Solving From Nature V - LNCS 1498*, pages 907–916. Springer-Verlag, Berlin, 1998.