

# Guided Genetic Algorithm and its Application to Radio Link Frequency Assignment Problems

T. L. Lau ([tllau@essex.ac.uk](mailto:tllau@essex.ac.uk))

*Dept. of Computer Science, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK*

E. P. K. Tsang ([edward@essex.ac.uk](mailto:edward@essex.ac.uk))

*Dept. of Computer Science, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK*

**Abstract.** The Guided Genetic Algorithm (GGA) is a hybrid of Genetic Algorithm and Guided Local Search, a meta-heuristic search algorithm. As the search progresses, GGA modifies both the fitness function and fitness template of candidate solutions based on feedback from constraints. The fitness template is then used to bias crossover and mutation. The Radio Link Frequency Assignment Problem (RLFAP) is a class of problem that has practical relevance to both military and civil applications. In this paper, we show how GGA can be applied to the RLFAP. We focus on an abstraction of a real life military application that involves the assigning of frequencies to radio links. GGA was tested on a set of eleven benchmark problems provided by the French military. This set of problems has been studied intensively by a number of prominent groups in Europe. It covers a variety of needs in military applications, including the satisfaction of constraints, finding optimal solutions that satisfy all the constraints and optimization of some objective functions whenever no solution exist ("partial constraint satisfaction"). Not only do these benchmark problems vary in problem nature, they are reasonably large for military applications (up to 916 variables, and up to 5548 constraints). This makes them a serious challenge to the generality, reliability as well as efficiency of algorithms. We show in this paper that GGA is capable of producing excellent results reliably in the whole set of benchmark problems.

**Keywords:** Genetic Algorithm, Constraint Satisfaction Optimization Problem, Partial Constraint Satisfaction Problem

## 1. Introduction

A finite Constraint Satisfaction Problem (CSP) can be described as a problem with a finite set of variables, where each variable is associated with a finite domain. Relationships between variables constrain the possible instantiations they can take at the same time [37, 15]. To solve a CSP, one must find the solution tuple that instantiate variables with values of their respective domains, and that these instantiations do not violate any of the constraints. Our area of research is in Constraint Satisfaction Optimization Problem (CSOP) and Partial Constraint Satisfaction Problem (PCSP), two variations of the CSP.



© 2000 Kluwer Academic Publishers. Printed in the Netherlands.

In the realms of CSP, the instantiation of a variable with a value from its domain is called a *label*. A simultaneous instantiation of a set of variables is called a *compound label*, which is a set of labels. A *complete compound label* is one that assigns values, from the respective domains, to all the variables in the CSP.

A CSOP is a CSP with an objective function  $f$  that maps every complete compound label to a numerical value. The goal is to find a complete compound label  $S$  such that  $f(S)$  gives an optimal value, and that no constraint is violated. A PCSP is similar to a CSOP except that the complete compound label may have variable instantiations that violate some of its constraints. Violation is unavoidable when the constraints are so tight that a satisfiable solution does not exist, or cannot be found [16, 37]. Deciding which constraint to violate is influenced by its cost and type. Hard constraints are types of constraints that must not be violated, whereas soft constraints may. The sum cost of all violated constraints is reflected in the objective function, further to other optimization criteria.

Guided Local Search (GLS) is a meta-heuristic algorithm that sits on top of hill-climbing algorithms to improve their efficiency and effectiveness. It has been applied to a number of problems and achieved outstanding results [21, 28, 38, 43, 41, 40, 44]. The objective of this research is to demonstrate that not only can GLS sit on top of hill-climbing algorithms, it can sit on top of GAs too. By sitting GLS on top of a GA, we get an algorithm which contains the epistasis in GA (through the use of GLS) and improves the robustness of GLS (contribution of GA). In this paper, we show how such a hybrid algorithm can be applied to the radio link frequency assignment problem. We claim that this work contributes to both meta-heuristic search and GA research.

## 1.1. THE RADIO LINK FREQUENCY ASSIGNMENT PROBLEM

### 1.1.1. *Background*

The Radio Link Frequency Assignment Problem (RLFAP) is a general problem that has relevance to both military and civil applications. RLFAP is NP-hard and is a variation on the T-graph colouring problem introduced in [18]. In this paper, we focus on the military-motivated frequency assignment problem tackled by the EUCLID CALMA (Combinatorial Algorithms for Military Applications) consortium. This was a group of six research bodies in Europe that was formed to investigate the use of combinatorial optimization techniques (including AI and OR techniques) to aid military decisions (the project ended in late 1995). It defined a class of RLFAPs that capture the needs in real life military applications. A set of benchmark problems was used to ob-

serve the effectiveness of different approaches. This set comprises eleven RLFAP instances with various optimization criteria. It was made publicly available through the efforts of the French Centre d'Electronique l'Armement.<sup>1</sup>

### 1.1.2. *Types of Instances*

Each instance in the CALMA benchmark has a set of files that describe its variables, their domains, the constraints, and the objective. In addition, we are also given information on the respective optimization requirements based on the solubility of the problem. Optimization criteria describe the interpretation of variable instantiations and the means of measuring their desirability; thus shaping the objective function of our search routine, whereas the solubility of the problem states if a solution can be found under the condition that no constraint was violated. If an instance can be solved without constraint violation, then its optimality is defined as either O1 or O2, otherwise it is O3 (see below). For an insoluble instance, we use the violation cost of each constraint to solve the instance as a PCSP. In this paper, regardless of the problem's solubility, all instances in RLFAP are solved as PCSP since a CSOP problem can be mapped into a PCSP by giving each constraint a violation cost. This violation cost is the same value for all constraints in the instance.

- O1 - optimal solution is one with the fewest number of different values in its variables.
- O2 - optimal solution is one where the largest global assigned value is minimal.
- O3 - if a the problem cannot be solved without violating constraints, find a solution that minimizes the objective function as follows:

$$\begin{aligned}
 & a_1 \times nc_1 + a_2 \times nc_2 + a_3 \times nc_3 + a_4 \times nc_4 + \\
 & + b_1 \times nv_1 + b_2 \times nv_2 + b_3 \times nv_3 + b_4 \times nv_4
 \end{aligned} \tag{1}$$

Where  $nc_i$  is the number of violated constraints of priority  $i$ ,  $nv_i$  is the number of modified variables with mobility  $i$ . Mobility for

---

<sup>1</sup> The RLFAP benchmark problems used in the CALMA project have up to 1,000 variables and 6,000 constraints. They are available at the Centre d'Electronique l'Armement (France), via ftp at ftp.cert.fr/pub/bourret. Civil RLFAPs are often bigger in size (up to 10,000 variables and 2.5 million constraints). Some of the military requirements may not apply to civil RLFAPs and vice versa. Therefore benchmark results in one set of problems may not be generalized to the other. For GA applications to civil RLFAPs, readers may consult [6, 7, 5, 9, 19, 10, 33].

a radio link states the cost for changing the frequency from its assigned default. The values of the weights  $a_i$  and  $b_i$  are given if necessary.

All constraints in the CALMA benchmark are binary; that is, each constraint operates on the values in two variables. These constraints test the absolute difference of two variables in a candidate solution, where this logical test can belong to either of the two following classes:

C1 - the absolute difference must be lesser than a constant.

C2 - the absolute difference must be equal to a constant.

Table I lists the instances, their characteristics and its objective. From this table, we can observe that the RLFAP contains instances that are varied in both the optimization and constraint criteria. Further, the number of variables, their domain sizes, and the number of constraints on these variables make the RLFAP a non-trivial problem set for any algorithm. The RLFAP would not only test the quality and robustness of an algorithm, but also its flexibility to adapt to the different optimization and constraint criteria of each instance.

Table I. Characteristics of RLFAP instances.

Instance	No. of variables	No. of constraints	Soluble	Minimize	Type
scen01	916	5548	Yes	Number of different values used	O1
scen02	200	1235	Yes	Number of different values used	O1
scen03	400	2760	Yes	Number of different values used	O1
scen04	680	3968	Yes	Number of different values used	O1
scen05	400	2598	Yes	Number of different values used	O1
scen06	200	1322	No	The maximum value used	O2
scen07	400	2865	No	Weighted constraint violations	O3
scen08	916	2744	No	Weighted constraint violations	O3
scen09	680	4103	No	Weighted constraint violations and mobility costs	O3
scen10	680	4103	No	Weighted constraint violations and mobility costs	O3
scen11	680	4103	Yes	Number of different values used	O1

## 2. The RLFAP in PCSP Expression

A PCSP is defined as a quadruple of  $\{Z, D, C, f\}$  where  $Z$  is a finite set of variables. With respect to  $Z$ ,  $D$  is a function that maps every variable to a set of values, which is called a domain.  $C$  is a finite set of constraints that affect a subset of the variables, and each constraint has a cost for its violation. The objective function  $f$  returns a magnitude based on the instantiation of the variables and the satisfaction of constraints. In the RLFAP, each instance has a set of files that conveniently describe the respective  $Z$ ,  $D$  and  $C$  sets, and the optimization objective  $f$ .

### 2.1. VARIABLES AND DOMAINS

For any of the RLFAP instance with  $m$  variables, let  $q_j$  be a variable in  $Z$  representing one radio link. For each variable  $q_j$  in  $Z$ , there is one associated domain mapped by the function  $D$ , denoted by  $D(q_j)$ , which contains a set of  $n$  values, each value representing a valid frequency that can be assigned to the variable.

$$Z = \{q_1, q_2, \dots, q_m\} \quad (2)$$

$$\text{where } \forall q_j \in Z : D(q_j) = \{freq_1, freq_2, \dots, freq_n\} \quad (3)$$

### 2.2. CONSTRAINTS

The constraint set  $C$  consists of  $n$  elements, representing  $n$  constraints in the instance. Each element in  $C$  consist of the constraint  $c_i$  and its cost  $cost_i$  (Eq. 4). In instances of type O3 (see section 1.1.2),  $cost_i$  holds an integer denoting the cost incurred if constraint  $c_i$  is violated. From Eq. 1, we see that in the RLFAP, each constraint has an assigned priority that is related to a user defined cost. We used this priority to derive the  $cost_i$  for each constraint  $c_i$ .

As discussed in section 1.1.2, there are two types of binary constraints in the RLFAP; C1 and C2 which we formulate into Eq. 5. In that equation,  $q_a$  and  $q_b$  are two variables from a candidate solution, and  $z$  is a constant. Eq. 6 states that constraint  $c_i$  returns a binary value that is 1 for a violation and 0 otherwise.

$$C = \{ \langle c_1, cost_1 \rangle, \langle c_2, cost_2 \rangle, \dots, \langle c_n, cost_n \rangle \} \quad (4)$$

$$\forall c_i \in C : \begin{cases} c_i \equiv |q_a - q_b| < z, \text{ if } c_i \text{ is type C1} \\ c_i \equiv |q_a - q_b| = z, \text{ if } c_i \text{ is type C2} \end{cases} \quad (5)$$

$$\text{where } c_i = \begin{cases} 1, \text{ constraint is violated} \\ 0, \text{ otherwise} \end{cases} \quad (6)$$

### 2.3. OBJECTIVE FUNCTION

The respective objective function  $f$  of the instances in RLFAP are stated in Table I. The objective functions are also explained in section 1.1.2.

## 3. Algorithms

CSPs and CSOPs are generally NP-hard [37] and although heuristics have been found useful in solving them, most systematic search algorithms are deterministic and constructive [27], and would thereby be limited by the combinatorial explosion problem. Systematic methods include search and inference techniques. These search methods are complete, so they are able to guarantee a solution, or to prove that one does not exist. Thus systematic techniques will, if necessary, search the entire problem space for the solution [22].

The combinatorial explosion is an obstacle faced by systematic search methods for solving realistic CSPs, and in looking for optimal and or near-optimal solutions in CSOPs. In optimization, to ensure that the solution found is the optimal, systematic search algorithms in the worst case would need to exhaust the entire problem space to establish that fact. Branch and bound methods help to prune off parts of the space. Their effectiveness depends on their ability to find good bounds, which is the subject of much research, e.g. see [32, 39, 3].

Stochastic search methods are normally incomplete. They are not able to guarantee that a solution can be found, and neither can they prove that a solution does not exist. They forgo completeness for efficiency. Often, stochastic search methods can be faster in solving CSOPs than systematic methods [14]. Many publications such as [29, 31, 23] demonstrated on several large problems that systematic search algorithms fail to solve, but stochastic alternatives efficiently conquer.

### 3.1. GENETIC ALGORITHMS

Genetic Algorithms (GAs) are stochastic search algorithms that borrow some concepts from nature [8, 17, 20]. GA maintains a *population pool*  $\gamma$  of candidate solutions called *strings* or *chromosomes*. Each chromosome  $\gamma_p$  is a collection of  $\alpha$  building blocks known as *genes*, which are instantiated with values from a finite domain. Let  $\gamma_{p,q}$  denote the value of gene  $q$  in chromosome  $p$  in the population  $\gamma$ .

Associated with each chromosome is a *fitness* value which is determined by a user defined function. The function returns a magnitude

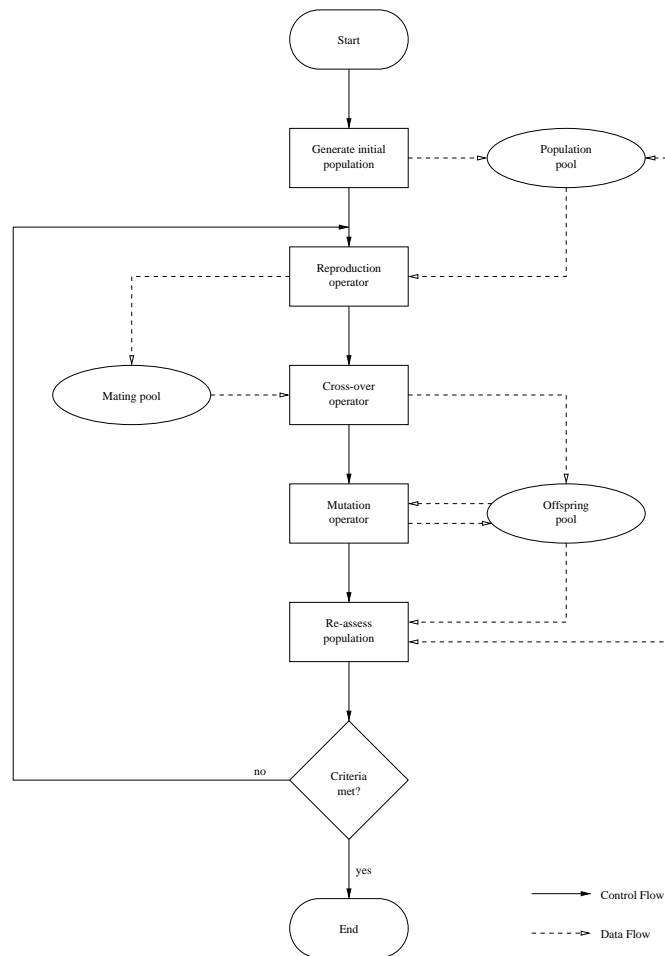


Figure 1. A canonical Genetic Algorithm

that is proportional to the candidate solution's suitability and/or optimality. Fig. 1 shows the control and data flow of a canonical GA. At the start of the algorithm, an initial population is generated. Initial members of the population may be randomly generated, or generated according to some rules. The *reproduction operator* selects chromosomes from the population to be parents for a new chromosome and enters them into the *mating pool*. Selection of a chromosome for parenthood can range from a totally random process to one that is biased by the chromosome's fitness.

The *crossover operator* oversees the mating process of two chromosomes. Two parent chromosomes are selected from the mating pool randomly and the *crossover rate*, which is a real number between zero and one, determines the probability of producing a new chromosome from the parents. If the mating was performed, a child chromosome is created which inherits complementing genetic material from its parents. The crossover operator decides what genetic material from each parent is passed onto the child chromosome. The new chromosome produced is entered into the *offspring pool*.

The *mutation operator* takes each chromosome in the offspring pool and randomly changes part of its genetic make-up, i.e. it's content. The probability of mutation occurring on any chromosome is determined by the user specified mutation rate. Chromosomes, mutated or otherwise, are returned to the offspring pool after the mutation process.

Thus each new generation of chromosomes are formed by the action of genetic operators (reproduction, crossover and mutation) on the older population. Finally, the members of the population pool are compared with those of the offspring pool. The chromosomes are compared via their fitness value to derive a new population, where the weaker chromosomes may be eliminated. In particular, weaker members in the population pool are replaced by the fitter child chromosomes from the offspring pool. The heuristic for assessing the survival of each chromosome into the next generation is called the *replacement strategy*.

The process of reproduction, crossover mutation and formation of a new population completes one *generation* cycle. A GA is left to progress through generations, until certain criteria (such as a fixed number of generations, or a time limit) are met. GAs were initially used for machine learning systems, but it was soon realised that GAs have great potential in function optimization [1, 17, 8].

### 3.2. MOTIVATION FOR RESEARCH

A basic assumption of the canonical GA is that genes are independent of each other, so that the value taken by one gene will not influence the instantiation of any other gene in the same chromosome. Constrained problems are highly epistatic in nature. Epistasis is the interaction between different genes in a chromosome. A candidate solution to a typical constrained problem is often represented as a chromosome, where each gene in the chromosome describes a variable in the problem. Constraints influence both the values that sets of genes can take simultaneously and the overall fitness of that chromosome. Goldberg suggested high epistasis as an explanation to GAs failure in certain



tasks [17]. To contain epistasis, many have developed specialized GAs for handling constraints, e.g. see [2, 11, 12, 13, 6, 7, 5, 30, 35, 47, 48, 46].

One way to contain the epistasis problem is to design a representation that will not generate illegal solutions through crossover. This is possible for some problems, but it may not be easy to find such representations for others. Epistasis has plagued GAs since its inception and the successful suppression of its effects would not only be significant to the current goal (that of using GAs for constrained problems), but also to the field of GAs. This will involve fundamental changes to the canonical GA architecture.

By maintaining a collection of solutions in parallel, and through the explorative mode (through the use of its crossover operator) of Genetic Algorithms, they are able to mine the search space in several locations at the same time. This makes Genetic Algorithms less local than compared to most local search techniques, giving them an edge of *robustness*. Robustness is defined as consistency in finding optimal (or near optimal) solutions, and the consistency in the quality of these solutions. This is a much desired quality that would boost confidence in the use of Artificial Intelligence in the real world. Since it would be difficult to entrust a mission critical problem to a method that gives solutions with a high degree of variance.

### 3.3. GUIDED GENETIC ALGORITHM

#### 3.3.1. *Background*

Among our earlier work on CSOPs, we looked at the Processor Configuration Problem (PCP) [4, 46]. Briefly, the PCP is a real life CSOP where the task is to link up a finite set of processors into a network, whilst minimizing the maximum distance between these processors. Since each processor has a limited number of communication channels, a carefully planned layout will help reduce the overhead for message switching.

We developed a GA called the Lower Power Genetic Algorithm (LPGA) [25, 26] specifically for solving the PCP. LPGA is a two-phase GA approach where in the first phase, we run LPGA until a local optimum has been determined. The best chromosome from this run is analysed and used to construct a fitness template for use in the next phase. This fitness template is a map that defines undesirable genes, so influencing LPGA to change their contents. By insisting that crucial genes do not change, the evolution in the second phase shifts focus onto other parts of the string; resulting in a more compact search space.

LPGA found improved solutions than results published so far in the PCP. It's success could be attributed to the use of an effective data

representation and more importantly, the presence of an application specific penalty algorithm. In our effort to generalize LPGGA, we sought to develop a GA that utilizes a dynamic fitness template constructed by a general penalty algorithm. The Guided Genetic Algorithm (GGA) reported in this paper was the result of this effort.

### 3.3.2. Overview of GGA

In our journey to develop GGA, we have taken liberty with some of the traditional GA concepts (such as the addition of a penalty operator, and an alternate interpretation of the mutation rate). These will be introduced as we progress through the rest of this paper. Comparing GGA in Fig. 2 against the canonical GA in Fig. 1, we see the addition of data collection called the fitness templates, a penalty operator (see 3.3.3) and a condition to activate that operator. Also added to Fig. 2 are the interactions between the penalty operator and the data space in GGA. Appended at the end of this section (section 3) are two tables (Tables II and III), summarizing for the readers' convenience, the terms and technology introduced henceforth.

The control flow of the GGA is very much similar to that of the canonical GA, described in section 3.1. After the start of the algorithm, an initial population is created. A new generation of chromosomes are derived from the parent chromosomes through the actions of the reproduction, crossover and mutation operators. Both the crossover and mutation operators (or any operator thereof) may be adapted to use the information provided by the penalty operator, via the *fitness template* of individual chromosome (explained in section 3.3.4). Memberships to the population pool are re-assessed by comparing the fitness of the chromosomes from the population and offspring pool. For the RLFAP, GGA was configured to use an *elitist replacement strategy*. Under this strategy, chromosomes from both the population and offspring pool are ranked by their fitness. The fittest  $n$  chromosomes in the ranking are used to form the next generation's population pool. In GGA,  $n$  is set to the size of the population pool.

New elements of the GGA come into play at this point. The new population is surveyed for the possibility of being trapped in a local optimum. We can observe that in hill climbing, when a search is trapped in a local optimum, it repeatedly returns the same solution since the neighbouring states does not offer any improvement. If the population is indeed trapped in a local optimum, the penalty operator is called. In GGA, we conclude that the best chromosome is trapped in a local optimum if its fitness does not improve over a given number of iterations. This number is a parameter to GGA. When the penalty operator is called, it looks for undesirable features in the chromosomes

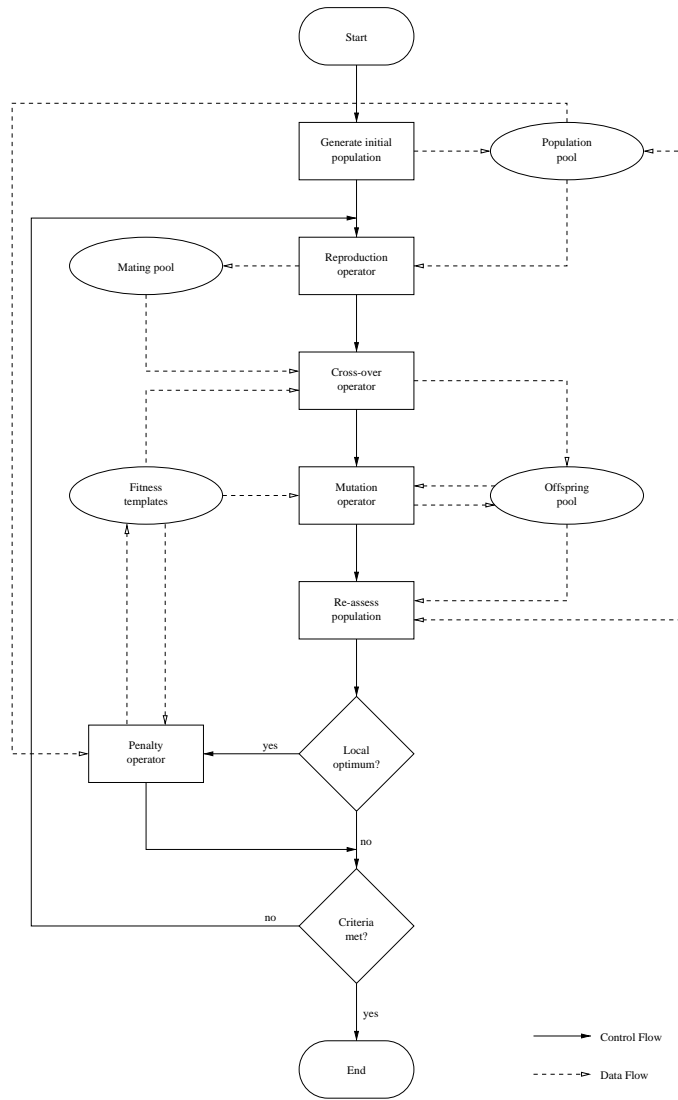


Figure 2. The Guided Genetic Algorithm

and update the fitness template (or fitness templates), so that mutation and crossover operators might fade out these features in the coming generations.

### 3.3.3. Penalty Operator

The use of a fitness template in LPGA (generated by a specialized penalty algorithm) was the motivating force in the development of GGA. In the quest for a general penalty algorithm, we looked for functional similarities to LPGA and more importantly, that the nature of the penalty algorithm will not be obstructive to the operation of a canonical GA. The Guided Local Search (GLS) [40] developed by our research group is an intelligent search scheme for combinatorial optimization problems. It meets our criteria and further, its conceptual simplicity and proven effectiveness in a range of well known problems was an added attraction [38, 42, 41, 43]. In GGA, we adapted GLS in the form of the *penalty operator*.

Solutions are characterized by a set of solution features  $\theta$ , where a solution feature  $\theta_i$  can be any property exhibited by the solution (Eq. 7). This property must be non-trivial, such that it does not appear in all candidate solutions. Research on GLS has indicated that, in many cases, feature definition is not difficult, since the domain often suggests features that one could use. The application in this paper supports this point.<sup>2</sup>

$$\theta = \{\theta_1, \theta_2, \dots, \theta_m\} \quad (7)$$

In GGA, a feature  $\theta_i$  is limited to variable assignments (in GLS, it is more general); a feature in a chromosome may be exhibited by the simultaneous assignments of a group of genes. Thus the feature  $\theta_i$  defines a set of positions in the chromosome representation, and is represented by an *indicator function*  $\tau_i$  in Eq. 8, which test the existence of that feature. For each feature  $\theta_i$ , there is a *cost*  $\eta_i$  which rates that feature's presence in a solution in degrees of undesirability. Indicator functions and costs are application dependent, and so they are defined by the user.

$$\tau_i(\gamma_p) = \begin{cases} 1, & \text{solution } \gamma_p \text{ exhibits feature } \theta_i \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

The penalty counter  $\zeta_i$  is a variable maintained by GGA that gives the degree of extent that the feature  $\theta_i$  is penalized as the search is progressing; the counter is initialized to zero at the beginning of the

---

<sup>2</sup> Features can often be defined from the objective function. For example, in the travelling salesman problem, whether one goes immediately from city  $x$  to city  $y$  can be defined as a feature[44]. This feature is only exhibited by tours that have cities  $x$  and  $y$  adjacent to each other. The cost of this feature is the distance to be travelled from  $x$  to  $y$ . We have been able to find features in other problems without difficulty. for example, see [24, 38, 40, 41, 44].

search and its value can only increase. A new fitness function called the *augmented cost function*  $g$  (Eq. 9) is used in place of function  $f$ , so that changes in penalty counters will affect the survival of chromosomes. The *regularization parameter*  $\lambda$  (adopted from GLS) measures the impact penalties have, with respect to function  $f$ .

$$g(\gamma_p) = f(\gamma_p) + \lambda \cdot \sum (\zeta_i \cdot \tau_i(\gamma_p)) \quad (9)$$

In GGA, if the fitness of the best chromosome remains unchanged for a specific number of generations, we conclude that it is trapped in a local optimum. The penalty operator comes in to analyze the best chromosome  $\gamma_p$  of the population for features to penalize. Penalties are used in GGA to guide the search to escape local optima. To evaluate the utility of penalizing individual features exhibited by a candidate solution, GGA (following GLS) takes into consideration the cost as well as the penalty counter (Eq. 10). Thus, for all features  $\theta_i$  in the fittest chromosome  $\gamma_p$  that maximizes the function  $util(\gamma_p, \theta_i)$  (Eq. 10), the related penalty counter  $\zeta_i$  is incremented by one. It is hoped that by penalizing undesirable features, we can escape from the local optimum and suppress the occurrence of these features in the coming generations.

$$util(\gamma_p, \theta_i) = \tau_i(s) \cdot \frac{\eta_i}{1 + \zeta_i} \quad (10)$$

#### 3.3.4. *Fitness Template*

Central to the theme in GGA is the fitness templates. Besides the fitness function, the fitness templates offer an added channel of communication between the penalty operator, and the mutation and crossover operators. The fitness template is a map that defines which genes in a chromosome are more susceptible to be changed during crossover or mutation.

In GGA, each chromosome  $\gamma_p$  in the population is associated with exactly one *fitness template*  $\delta_p$ . Each fitness template is made up of smaller units known as *weights*  $\delta_{p,q}$ , each of which corresponds to a gene  $\gamma_{p,q}$ . A weight  $\delta_{p,q}$  is a positive integer. The “heavier” a gene appears (compared to its comrades), the greater are its chances of having its content altered. Therefore in the case of mutation, the weight of a gene is proportional to the probability that mutation may occur on it, relative to the weights of other genes in the same chromosome. This is especially useful when the number of genes in a chromosome is large, where random selection of genes might not be helpful. More details on the role of the fitness template with the mutation and crossover operator will be given in their respective sections.

Weights in the fitness template for each chromosome are computed when the chromosome was first created, and after the penalty operator has penalized feature(s). Computation of weights are needed after these events because the content of either the chromosomes or the penalty counters have changed.

For a chromosome, the distribution process (Fig. 3) starts by initializing all weights to zero. It will check the chromosome for the presence of any features from the set  $\theta$ . For a feature  $\theta_i$  that exist in the chromosome<sup>3</sup>, all the weights related to the gene positions defined by  $\theta_i$  is incremented with the value in its penalty counter  $\zeta_i$ .

```

FUNCTION DistributePenalty( chromosome  $\gamma_p$  )
{
  FOR EACH weight  $\delta_{p,q}$  RELATED TO chromosome  $\gamma_p$ 
  {
     $\delta_{p,q} \leftarrow 0$ 
  }

  FOR EACH solution feature  $\theta_i$  IN feature set  $\theta$ 
  {
    IF  $\tau_i(\gamma_p) = 1$  THEN
    {
      FOR EACH gene position  $q$  defined by  $\theta_i$ 
      {
         $\delta_{p,q} \leftarrow \delta_{p,q} + \zeta_i$ 
      }
    }
  }
}

```

Figure 3. Algorithm for the Distribution of Penalty

### 3.3.5. Cross-over Operator

In GGA, the action of mating two individuals from the population produces a new child. Each parent contributes a set of genes which the child inherits. In GA, the process of choosing parents, deciding their respective contribution rights of genetic material, and the forging of a child chromosome from these material is the responsibility of the crossover operator. The probability of crossover occurring is controlled

<sup>3</sup> Feature  $\theta_i$  is present when its indicator function  $\tau_i$  returns a one.

by the parameter *crossover rate*. By assembling a new chromosome that contains parts of two parent chromosomes, it may introduce to the population a new point in the search space. And since the parents chosen for mating are selected with bias to their fitness, we hope that the child chromosome may be fitter.

Cross-over operators differ primarily from each other in the way that they choose the genes from the parents to form the child. In the canonical GA, one of the simplest form of crossover is the *one-point crossover* [17, 8]. In Fig. 4, we have two parent chromosomes whose genes are binary encoded. One random point along the length of the chromosomes are selected as the crossover point. Each parent donates one different part of their chromosome (defined by the crossover point) to create the child chromosome.

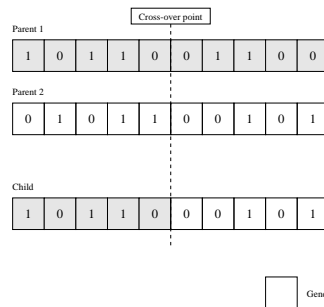


Figure 4. An example of the One-point Cross-over Operator in action

In GGA, we have adapted the crossover operator to take advantage of the fitness template. Two chromosomes  $\gamma_p$  and  $\gamma_{p'}$  are selected as parents to produce the child  $\gamma_{p''}$ . Each gene in chromosome  $\gamma_p$  competes against the corresponding gene in  $\gamma_{p'}$  for a place in the child. This competition is a weighted random selection, influenced by the weights  $\delta_{p,q}$  and  $\delta_{p',q}$  of the respective genes; thus the “lighter” gene will have a greater chance to propagate its information to the child. Note that the child does not inherit the weights from its respective parent, since the child may represent a different solution from its parents, and thus requiring the penalty operator to re-assess it. The algorithm of GGA’s crossover operator is shown in Fig. 5. Fig. 6 shows its effect when applied to the situation for one-point crossover in Fig. 4.

The operator starts by receiving two parents  $\gamma_p$  and  $\gamma_{p'}$  from the mating pool. For each set of corresponding genes  $\gamma_{p,q}$  and  $\gamma_{p',q}$  in the parents, it computes the *sum* of their weights. The selection of the gene is randomly biased, such that the probability for either  $\gamma_{p,q}$  or  $\gamma_{p',q}$

to have its gene passed onto their child is  $\frac{\gamma_{p',q}}{sum}$  and  $\frac{\gamma_{p,q}}{sum}$  respectively; giving the advantage to a “lighter” gene, which we wish the child  $\gamma_{p'',q}$  to inherit. This gene selection process is repeated for all genes in the parents. When a child chromosome is complete, its fitness and weights are computed.

```

FUNCTION CrossOver( parent chromosomes  $\gamma_p$  and  $\gamma_{p'}$  )
{
  FOR EACH gene position  $q$  IN the chromosome
  {
     $sum \leftarrow \delta_{p,q} + \delta_{p',q}$ 
     $point \leftarrow$  random integer from  $\{0, \dots, sum - 1\}$ 

    IF  $point < \delta_{p,q}$  THEN
    {
       $gene \leftarrow \gamma_{p',q}$ 
    }
    ELSE
    {
       $gene \leftarrow \gamma_{p,q}$ 
    }
  }

  RETURN gene as  $\gamma_{p'',q}$  for the offspring
}

```

Figure 5. Algorithm of the GGA Cross-over Operator

### 3.3.6. Mutation Operator

Mutation produces variations in the population through altering the information that genes carry. The *mutation rate* states the probability that mutation may occur on a chromosome. In GGA, the mutation rate is defined as a fraction of the size of each chromosome; the number of genes in a chromosome to mutate is the product of the mutation rate and the size of that chromosome.

In GGA, mutation (Fig. 7) acts on every child chromosome  $\gamma_{p''}$  produced by the crossover operator. For each chromosome, a number of genes are chosen (as above, decided by the mutation rate) to be modified. A gene  $\gamma_{p'',q}$  is selected using the *roulette wheel selection method*. In this selection method, the probability for each gene to be picked



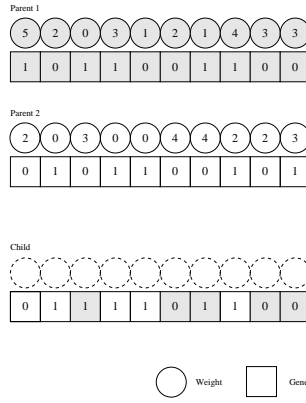


Figure 6. An example of the GGA Cross-over Operator in action (Weights for the child is calculated afresh, not inherited)

is directly proportional to its weight. Thus a gene with a “heavier” weight (and therefore less desirable) compared to others in the chromosome, will have a greater chance of being selected. Appended below is a description of our implementation of the roulette wheel selection method.

Given the chromosome  $\gamma_{p''}$ , we compute the sum of all the weights in the fitness template associated to this chromosome as  $sum = \sum_{q=1}^{\alpha} \delta_{p'',q}$ . The probability that a gene  $\gamma_{p'',q}$  is selected is proportional to its weight over  $sum$ , i.e.  $P(\gamma_{p'',q}) = \frac{\delta_{p'',q}}{sum}$ .

The next step for a selected gene  $\gamma_{p'',q}$  is to seek an appropriate value for replacement. This could be totally random or in GGA’s case, a value that will derive the best fitness (the biggest improvement) for the chromosome. In our algorithm, we have the variables *best* and *list*. The variable *best* holds the best fitness value, while *list* contains a list of values that will allow the chromosome to arrive at the fitness value in *best*. We step through all the values  $x_j$  in the domain  $D(q)$  relevant to the gene  $\gamma_{p'',q}$ . If  $x_j$  produces a new fitness  $z$  greater<sup>4</sup> than *best*, *best* is set to  $z$  and *list* is emptied. However, if the  $z$  is equal to *best*, the value  $x_j$  is added to the *list*. When all values in the domain have been exhausted, we randomly instantiate  $\gamma_{p'',q}$  with a value from *list*. Since at this point, *list* should contain all possible values that will give the chromosome  $\gamma_{p''}$  the biggest improvement.

<sup>4</sup> Since RLFAP is a minimization problem, we would want the greatest descent.

```

FUNCTION Mutation( chromosome  $\gamma_{p''}$  )
{
   $i \leftarrow 0$ 

  WHILE  $i < \text{mutation rate} \times \alpha$  (length of chromosome)
  {
     $q \leftarrow \text{RouletteWheel}(\gamma_{p''})$ 
     $best \leftarrow g(\gamma_{p''})$ 
     $list \leftarrow \gamma_{p'',q}$ 

    FOR EACH value  $x_j$  IN domain  $D(q)$ 
    {
       $\gamma_{p'',q} \leftarrow x_j$ 
       $z \leftarrow g(\gamma_{p''})$ 

      IF  $z \geq best$  THEN
      {
        IF  $z > best$  THEN
        {
           $best \leftarrow z$ 
           $list \leftarrow \{\}$ 
        }
      }

       $list \leftarrow list + x_j$ 
    }

     $\gamma_{p'',q} \leftarrow \text{random value in } list$ 
     $i \leftarrow i + 1$ 
  }

  RETURN the mutated chromosome  $\gamma_{p''}$ 
}

```

Figure 7. Algorithm of the GGA Mutation Operator

Updating of a gene's weight takes place after it's value has changed, where the weight associated with it is reduced by one unit <sup>5</sup> so that the probability of the same gene getting selected by the roulette wheel selector is reduced.

Gene mutation is repeated until the stopping criteria is met. As stated before, we stop mutating when the number of genes changed have reached a value that is the product of the mutation rate and the chromosome's length.

#### 4. Preparing GGA to solve RLFAP

In section 2, we expressed the RLFAP as a formal PCSP. In this section, we discuss the steps needed to adapt those definitions into a form that GGA can use.

The feature set  $\theta$  is a union of the feature set of constraints  $\theta_{cst}$  and the set of mobility of radio links  $\theta_{mbt}$  (Eq. 11). Constraint  $c_i$  defined in Eq. 5 is recast as a feature in the set  $\theta_{cst}$  (Eq. 12), where a one is returned if the constraint cannot be satisfied, and zero otherwise. The value of cost  $\zeta_{cst_i}$  to each constraint  $c_i$  depends on the nature of the instance. If the instance is soluble, then all  $\zeta_{cst_i}$  are set to a large value; say 10000, to signify that the constraint must not be broken (i.e. hard constraints). For insoluble instances,  $\zeta_{cst_i}$  is set to the weights given for its priority class (see section 1.1.2). Similar to soluble instances, hard constraints in insoluble instances will have their  $\zeta_{cst_i}$  set to a large value. The set  $\theta_{cst}$  has  $n$  features, where  $n$  is the number of constraints in the instance.

For the O3 objective type of instances, we need to minimize the mobility cost of our candidate solution, in addition to minimizing constraints violation costs. The set of mobility cost defines our next feature set,  $\theta_{mbt}$  (Eq. 13). For each variable in these instances, there is a mobility cost  $\zeta_{mbt_i}$  and a default assigned frequency  $default_i$ . If in our candidate solution, a variable has been assigned a value different from its default  $default_i$ , then a one is returned and zero otherwise. The mobility cost  $\zeta_{mbt_i}$  is set to the weights given for its priority class (again see section 1.1.2). There are radio links whose frequency should never change, and the mobility cost for these have been set to a large value. The feature set  $\theta_{mbt}$  defines  $n$  features, where  $n$  is the number of variables in the instance.

---

<sup>5</sup> Since a weight is a positive integer, a weight will only be decremented if it is greater than zero. By design, the same variable is allowed to mutate more than once because after  $x$  is mutated, mutation of other variables may result in  $x$ 's current value being locally suboptimal.

Table II. Components of GGA

Algorithms	Purpose
Cross-over operator	Uses the fitness templates of two parent chromosomes to decide each parent's contribution of genetic material towards creating a child chromosome.
Mutation operator	The fitness template of a chromosome is used to guide in the alteration of the chromosome's genetic content.
Penalty operator	This operator detects and selects undesirable solution features in a chromosome to penalize. Penalization involves incrementing penalty counters of the associated features.
Local optimum detector	Detects if the search is trapped in a local optimum. If it is, the penalty operator is called.
Distribute Penalty	If a solution feature is present in a chromosome, the penalty counter associated with this feature is added onto the weights of the genes that are constituents of this feature.
Data structures	Purpose
Weight $\delta$	Each gene has one weight. The weight is a measure of undesirability of the gene's current instantiation, compared to the rest of the chromosome.
Fitness template	A fitness template is a collection of weights. Each chromosome has one fitness template.
Solution feature $\theta$	Solution features are domain specific and user defined. A feature is exhibited by a set of variable assignments that describes a non-trivial property of a problem.
Penalty counter $\zeta$	Each feature has one penalty counter. A penalty counter keeps count of the number of times its related solution feature has been penalized since the start of the search.

Table III. Inputs and Parameters to GGA

Inputs/Parameters	Purpose
Solution feature $\theta$	See Table II
Cost $\eta$	Each solution feature has a cost to rate its undesirability of presence.
Indicator function $\tau$	Each solution feature has a user defined indicator function that tests for the feature's presence in a chromosome.
Objective function $f$	A function that maps each solution to a numerical value.
Regularization parameter $\lambda$	A parameter that determines the proportion of contribution that penalties have in an augmented fitness function.
Augmented fitness function $g$	A function that is the sum of the objective function $f$ on a chromosome and the penalties of features that exist in it. This is the function that GGA uses to compute fitness for each chromosome.
Mutation rate	A fraction that defines the number of genes in the chromosome to mutate.
Cross-over rate	The probability that crossover will occur between two chromosomes.

$$\theta = \{\theta_{cst}, \theta_{mbt}\} \quad (11)$$

$$\forall \theta_i \in \theta_{cst} : \tau_{cst_i}(\gamma_p) \equiv \begin{cases} 1, \text{ if C1 and } |\gamma_{p,a} - \gamma_{p,b}| \geq z_i \\ 1, \text{ if C2 and } |\gamma_{p,a} - \gamma_{p,b}| \neq z_i \\ 0, \text{ otherwise} \end{cases} \quad (12)$$

$$\forall \theta_i \in \theta_{mbt} : \tau_{mbt_i}(\gamma_p) \equiv \begin{cases} 1, \text{ if } \gamma_{p,i} \neq default_i \\ 0, \text{ otherwise} \end{cases} \quad (13)$$

For all instances in the RLFAP, we seek to minimize the function  $g$  (Eq. 14). In  $g$ , the function  $f$  depends on the objective type of the instance (Eq. 15). The cost  $\zeta_i$  and  $\tau_i$  both refers to the unified feature set of  $\theta$ . They will associate with  $\zeta_{cst_i}$  and  $\tau_{cst_i}$ , or  $\zeta_{mbt_i}$  and  $\tau_{mbt_i}$  where applicable. Thus the value  $n + c$  in Eq. 14 is the sum of the number of features in  $\theta_{mbt}$  (which defines  $n$  features, where  $n$  is the number of

variables) and  $\theta_{cst}$  (which defines  $c$  features, where  $c$  is the number of constraints).

$$g(\gamma_p) = f(\gamma_p) + \lambda \cdot \sum_{i=1}^{n+c} (\zeta_i \cdot \tau_i(\gamma_p)) \quad (14)$$

$$f(\gamma_p) = \begin{cases} \text{if O1, number of different values used in } \gamma_p \\ \text{if O2, largest value used in } \gamma_p \\ \text{if O3, } a_1 \times nc_1 + a_2 \times nc_2 + a_3 \times nc_3 + a_4 \times nc_4 + \\ \quad + b_1 \times nv_1 + b_2 \times nv_2 + b_3 \times nv_3 + b_4 \times nv_4 \end{cases} \quad (15)$$

## 5. Benchmark

The CALMA benchmark results by algorithms devised within the CALMA group were reported by Tiourine et al. in [36]. In this section, we compare GGA's results with the CALMA algorithms (section 5.2.1). Further, we will also evaluate the examine the value that GGA adds to the canonical GLS (section 5.2.3).

### 5.1. TEST ENVIRONMENT

In our physical environment, GGA was written in C++ and compiled using GNU GCC version 2.7.1.2. The code runs on an IBM PC compatible with a Pentium 133 MHz processor, 32MB of RAM and 512KB of Level 2 cache. Both compilation and execution of GGA was performed on the Linux operating system, using kernel version 2.0.27. Under GGA's environment, we have a mutation rate and crossover rate of 1.0, a population size of 20, and a stopping criterion of 100 generations. For the fitness function  $g$ ,  $\lambda$  has a value of 10.

### 5.2. PERFORMANCE EVALUATION

#### 5.2.1. Comparing all Algorithms

In Table IV, we see the published results of all the CALMA algorithms, GLS and GGA. Algorithms from the CALMA project groups consist of either complete or stochastic methods. The results recorded in the table are from the best solution each algorithm had generated. For soluble instances (scen01, scen02, scen03, scen04, scen05 and scen11)<sup>6</sup>, we report the number of frequencies above the known optimum that

<sup>6</sup> Instance scen06 was found to be insoluble, and thus solved as an O3 problem.

each solution (generated by the respective algorithms) uses <sup>7</sup>. Results for the insoluble instances are reported as the percentage deviation from the best known reported solution. When available, the average time taken for each algorithm to arrive at these solutions are also reported in the table.

For soluble instances, we observe that only seven out of the thirteen algorithms<sup>8</sup> were able to provide a solution to all the instances. Of the six soluble instances, GGA failed to return an optimum solution only for instance *scen11*. This instance has proved to be difficult for most of the algorithms, since only three algorithms (Taboo Search (EUT), Branch and Cut (DUT,EUT) and Constraint Satisfaction (LU)) were able to return a solution on par with the best known. Only two algorithms (Branch and Cut (DUT,EUT) and Constraint Satisfaction (LU)) were able to report solutions that gave the most optimal result. However, these two algorithms were limited to solving soluble instances only.

Looking at insoluble instances, we see that eleven out of the fourteen algorithms were able to tackle these PCSP instances. Of the eleven, nine of these algorithms managed to provide a satisfactory solution to the instances. Genetic Algorithms (LU) found the best known solutions ("in hours").

Overall, we see that top performers in each category (soluble and insoluble) are limited (in application) to only that category; Branch and Cut (DUT,EUT) and Constraint Satisfaction (LU) for soluble instances, and Genetic Algorithms (LU) for insoluble instances. Of the total of fourteen algorithms, ten were applicable to both categories. Of these ten, only four algorithms were able to find solutions to all the eleven instances. Of these four, Simulated Annealing (EUT) found good solutions in all instances except in *scen07*, in which it found a solution whose cost was 65% above the best known solution. Tabu Search (KCL) found good solutions in the soluble problems, but poorer solutions in insoluble instances, especially *scen06*, *scen07* and *scen08*, in which it found solutions whose costs were over 100% above the best known cost. GLS and GGA found solutions whose costs were within 10% of the best known costs in each of the eleven instances. In many cases, it found solutions as good as the best known solutions. This shows the consistency of GLS and GGA.

---

<sup>7</sup> Since the preparation of this paper, new results have been reported. Voudouris and Tsang limited the maximum number of penalties allowed in GLS (in a tabu-list manner) and obtained better results [45].

<sup>8</sup> The Genetic Algorithms from (LU) was designed for PCSP, and so it did not attempt any of the soluble instances.

Table IV. Comparison of GGA with GLS and the CALMA project algorithms.

Instance (scen)	Soluble Instances						Time	Platform
	01	02	03	04	05	11		
Simulated Annealing (EUT)	2	0	2	0	0	2	1min	SUN Sparc 4
Taboo Search (EUT)	2	0	2	0	-	0	5min	SUN Sparc 4
Variable Depth Search (EUT)	2	0	2	0	-	10	6min	SUN Sparc 4
Simulated Annealing (CERT)	4	0	0	0	-	10	41min	SUN Sparc 10
Tabu Search (KCL)	2	0	0	0	0	2	40min	DEC Alpha
Extended GENET (KCL)	0	0	0	0	0	2	2min	DEC Alpha
Genetic Algorithms (UEA)	6	0	2	0	-	10	24min	DEC Alpha
Genetic Algorithms (LU)	-	-	-	-	-	-	-	DEC Alpha
Partial Constr. Satisf. (CERT)	4	0	6	0	0	-	28min	SUN Sparc 10
Potential Reduction (DUT)	0	0	2	0	0	-	3min	HP 9000/72
Branch and Cut (DUT,EUT)	0	0	0	0	0	0	<10min	-
Constraint Satisfaction (LU)	0	0	2	0	0	0	hours	PC
Guided Local Search (UE)	0	0	0	0	0	6	20sec	DEC Alpha
<b>Guided Genetic Alg. (UE)</b>	0	0	0	0	0	2	40min	PC Linux

Best known solution      16   14   14   46   792   22

Results for the soluble instances are reported as the number of frequencies more than the optimum used.

Instance (scen)	Insoluble Instances					Time
	06	07	08	09	10	
Simulated Annealing (EUT)	6%	65%	5%	0%	0%	310min
Taboo Search (EUT)	-	-	-	-	-	-
Variable Depth Search (EUT)	3%	0%	14%	0%	0%	85min
Simulated Annealing (CERT)	42%	1299%	70%	2%	0%	42min
Tabu Search (KCL)	167%	1804%	566%	8%	1%	111min
Extended GENET (KCL)	12%	27%	40%	-	-	20min
Genetic Algorithms (UEA)	0%	386%	134%	3%	0%	120min
Genetic Algorithms (LU)	0%	0%	0%	0%	0%	hours
Partial Constr. Satisf. (CERT)	83%	2563%	246%	47%	12%	6min
Potential Reduction (DUT)	27%	-	-	4%	1%	10min
Branch and Cut (DUT,EUT)	-	-	-	-	-	-
Constraint Satisfaction (LU)	-	-	-	-	-	-
Guided Local Search (UE)	4%	9%	7%	0.7%	0.003%	2.88min
<b>Guided Genetic Alg. (UE)</b>	4%	9%	7%	0.7%	0.003%	60min

Best known solution      3437   343594   262   15571   31516

Results for the insoluble instances are reported as the percentage deviation from the best known solution (in cost, according to the objective functions specified in the problem). The best known solutions were found empirically, hence not necessarily optimal.

CERT    Centre d'Etudes et de Recherces de Toulouse, France  
DUT    Delft University of Technology, The Netherlands  
EUT    Eindhoven University of Technology, The Netherlands  
KCL    King's College London, United Kingdom  
LU    Limburg University, Maastricht, The Netherlands  
UEA    University of East Anglia, Norwich, United Kingdom  
UE    University of Essex, United Kingdom



### 5.2.2. Comparing Genetic Algorithms

Genetic Algorithms (UEA) [34], Genetic Algorithms (LU) [34] and GGA are the three GAs reported here. Genetic Algorithms (UEA) are a set of GAs based on the canonical GA but with specialized genetic operators, and are able to tackle both soluble and insoluble instances. The GGA performed better than Genetic Algorithms (UEA) in all but two instances (scen06 and scen10). Note that GGA had a clear edge over solutions from Genetic Algorithms (UEA) for instances scen01, scen07, scen08 and scen11.

Genetic Algorithms (LU) has genetic operators that exploit domain knowledge. This GA type algorithm has shown to be very effective for insoluble instances of the CALMA benchmark, as it was able to return solutions that are currently the best known. However, this achievement comes at the price of expensive computation, as Genetic Algorithms (LU) requires a large population size and long (hours of) processing time (when compared to Genetic Algorithms (UEA) and GGA).

### 5.2.3. Comparing GGA and GLS

Evaluating an algorithm's potential should not be limited to just the quality of the best solution it recommends. Besides solution quality and computation speed, one other measure is robustness. Robustness of an algorithm measures the consistency of the solutions it returns. In certain environments, one may need to be reasonably sure that the solution returned by an algorithm is, or very near the optimum. Unfortunately, we do not have any information on the robustness of the algorithms in the CALMA project. Therefore, we compare the statistical results of GGA and GLS.

For the comparison, GGA runs with a population of only five chromosomes. This is the minimum number of chromosomes needed to maintain an advantage over GLS, and the minimum population size to achieve acceptable results. Increasing the population size will improve robustness of GGA, but with diminishing significance. We introduce a variation of GLS, called *GLS5* to compete with GGA on even grounds. *GLS5* is five GLS running concurrently of each other, each maintaining its own candidate solution. Run time and iteration count for each GLS thread within *GLS5* has also been extended to meet with GGA's. That is, each GLS thread in *GLS5* has a limit of 100 iterations to match the 100 generations limit for GGA. At the end of each run, only the best solution from *GLS5* was used. The three algorithms (GGA, GLS and *GLS5*) were each executed 50 times, for each of the eleven instances.

The results from Table V shows GGA to have better robustness than GLS or *GLS5* in soluble instances, and in the case of scen11, GGA had reached a better solution. Also, GGA practically guarantees finding

Table V. Comparing Robustness between GGA, GLS and GLS5.

Inst.	Best Kn. Solution	Best Cost			Average Cost			Standard Deviation		
		GLS	GLS5	GGA	GLS	GLS5	GGA	GLS	GLS5	GGA
scen01	16	16	16	16	18.6	17.0	16.0	2.3	0.8	0.0
scen02	14	14	14	14	14.0	14.0	14.0	0.0	0.0	0.0
scen03	14	14	14	14	15.4	14.4	14.0	1.3	0.4	0.0
scen04	46	46	46	46	46.0	46.0	46.0	0.0	0.0	0.0
scen05	792	792	792	792	792.0	792.0	792.0	0.0	0.0	0.0
scen11*	22	28	26	22	-	33.9	30.2	-	3.1	1.7
scen06	3437	3575	3575	3575	4333.8	4129.6	4051.6	766.0	538.2	529.3
scen07	343594	374517	374517	374517	530641.1	510532.5	513044.8	79666.7	75149.4	75612.8
scen08	262	280	280	280	335.7	322.6	320.5	34.7	23.1	21.5
scen09	15571	15680	15680	15680	15999.7	15895.0	15889.6	194.7	112.6	109.3
scen10	31516	31517	31517	31517	31686.6	31631.4	31626.1	146.1	108.7	102.8

\* For *scen11*, results for GLS could not be computed because it did not return a solution that satisfied all constraints for some runs.

the optimum solutions for soluble instances scen01, scen02, scen03, scen04 and scen05. Results for insoluble instances are mixed, with both algorithms having results very close to each other. But overall, GGA achieved better average cost and smaller standard deviations for four of the five insoluble instances (over GLS5).

In Table IV, we see that solution quality reported by GLS and GGA were very much the same except in soluble instance *scen11*, where GGA managed to better GLS's result. The amount of CPU time required for computing these results have shown GLS to be much superior. However, it is unfair to view GGA as just a parallel version of GLS. In section 3.3, we describe the mechanics of GGA. We have integrated GLS as a component of GGA; i.e. as the penalty operator. The feedback from the GLS is used at two levels within GGA. On one level, GLS modifies the objective (fitness) function of GGA to influence its search. On another level, information from GLS is encoded into the template of each chromosome, which rates the relative fitness of each gene in the chromosome. The templates are used to influence the crossover and mutation processes. However, since GGA manages several candidate solutions at the same time, it will have a more complex means to detect local optimum traps.

## 6. Conclusion

We have demonstrated how GLS, as a meta-heuristic, can be sat on top of GA. We have designed a new meta-heuristic guided GA, which we call GGA, for arresting the effect of high epistasis when GAs are

deployed to solve problems such as those from the CSP class. On the other hand, the GA in GGA improves the robustness of GLS.

In this paper, we have explained how GGA could be applied to a particular set of RLFAPs. In a set of extensively studied benchmarks, we have shown that GGA adds value to the canonical GLS. And that overall, GGA performed well against the other algorithms. It is worth re-iterating that the set of benchmark problems cover a variety of needs and has been tackled seriously by a number of prominent research groups. GGA was able to produce results comparable to, and in many cases better than, these other methods in all the benchmark instances. This demonstrates the effectiveness and consistency of GGA.

The integration of GLS and the introduction of new elements to the foundation of the canonical GA gave GGA a technique of approximating gene fitness for a chromosome, and the provision for multi-criteria optimization. By knowing a gene's fitness within a chromosome, one could understand the magnitude of its contribution to the overall fitness. Gene fitness influence the effects that genetic operators have on them, encouraging change to genes with low fitness, whilst protecting the healthy ones.

For most applications where the users are more concerned with turnaround time and less so on robustness, GLS is clearly the better choice. But for mission critical applications, or applications where time is not as tight, GGA offers robustness and may perform better than GLS at times.

### Acknowledgement

The authors would like to thank the anonymous referees for their thorough and constructive comments. Chris Voudouris provided us with valuable advice on this work. Tung Leng Lau was sponsored by the University of Essex Studentship. Mathias Kern kindly helped us with formatting.

### References

1. Bethke, A.: 1978, 'Genetic Algorithms as Function Optimizers'. Technical Report 197, Logic of Computer Group, University of Michigan, USA.
2. Bowen, J. and G. Dozier: 1995, 'Solving Constraint Satisfaction Problems using a Genetic/Systematic Search Hybrid that Realizes When to Quit'. In: *Proceedings, 6th International Conference on Genetic Algorithms*. 122-129.
3. Cabon, B., S. de Givry, and G. Verfaile: 1996, 'Anytime Lower Bounds for Constraint Violation Minimization Problems'. In: *Proceedings, 4th International Conference on Principles and Practice of Constraint Programming*. 117-131.

4. Chalmers, A. and S. Gregory: 1993, 'Constructing minimum path configurations for multiprocessor systems'. *Parallel Computing* **19**, 343–355.
5. Chu, P.: 1997, 'Genetic Algorithms for Combinatorial Optimization Problems'. Ph.D. thesis, The Management School, Imperial College, University of London, UK.
6. Chu, P. and J. E. Beasley: 1996, 'Genetic Algorithms for the Generalized Assignment Problem'. *Computers and Operations Research*.
7. Chu, P. and J. E. Beasley: 1997, 'Constraint handling in genetic algorithms: the set partitioning problem'. Technical report, The Management School, Imperial College, University of London, UK.
8. Davis, L.: 1991, *Handbook of Genetic Algorithm*. Von Nostrand Reinhold.
9. Dorne, R. and J. K. Hao: 1995, 'An evolutionary approach for Frequency Allocation Problem in cellular radio-networks'. In: *IEEE International Conference on Evolutionary Computing (ICEC'95)*.
10. Dorne, R. and J. K. Hao: 1996, 'Constraint Handling in Evolutionary Search: A Case Study of the Frequency Assignment'. In: *Springer Verlag Lecture Notes in Computer Science vol. 1141, Int. Conf. Parallel Problem Solving from Nature (PPSN IV)*. 801–810.
11. Eiben, A., P.-E. Raué, and Z. Ruttkay: 1993, 'Heuristic genetic algorithms for constrained problems'. In: *Proceedings of Dutch National AI Conference NAIC '93*. pp. 241–252.
12. Eiben, A., P.-E. Raué, and Z. Ruttkay: 1994, 'Solving constraint satisfaction problem using genetic algorithms'. In: *Proceedings of 1st IEEE World Conference on Computational Intelligence*. 543–547.
13. Eiben, A., P.-E. Raué, and Z. Ruttkay: 1995, 'GA-easy and ga-hard constraint satisfaction problems'. In: *Proceedings of the ECAI-94 workshop on Constraint Processing*. 267–284.
14. Freuder, E., R. Dechter, B. Selman, M. Ginsberg, and E. Tsang: 1995, 'Systematic Versus Stochastic Constraint Satisfaction'. In: *Proceedings of 14th International Joint Conference on Artificial Intelligence*.
15. Freuder, E. and A. Mackworth: 1994, *Constraints-based Reasoning*. MIT Press.
16. Freuder, E. and R. Wallace: 1992, 'Partial constraint satisfaction'. *Artificial Intelligence* **58**, 21–70.
17. Goldberg, D. E.: 1989, *Genetic Algorithm in Search, Optimization and Machine Learning*. Addison-Wesley Pub. Co., Inc.
18. Hale, W. K.: 1980, 'Frequency assignment: Theory and applications'. In: *Proceedings of the IEEE*, Vol. 68. 1497–1514.
19. Hao, J. K. and R. Dorne: 1995, 'Study of Genetic Search for the Frequency Assignment Problem'. In: *Springer Verlag Lecture Notes in Computer Science vol. 1063 Artificial Evolution (AE'95), Brest, France*.
20. Holland, J.: 1965, 'Some practical aspects of Adaptive Systems Theory'. *Electronic Information Handling* 209 – 217.
21. Kilby, P., P. Prosser and P. Shaw: 1997, 'Guided local search for the vehicle routing problem'. *Proc., 2nd International Conference on Metaheuristics (MIC97)*. Sophia-Antipolis, France, pp.21-24
22. Kumar: 1992, 'Algorithms for Constraint Satisfaction Problems: A Survey'. *AI Magazine* **13**(1), 32–44.
23. Langley, P.: 1992, 'Systematic and non-systematic search strategies'. In: *Proceedings of Artificial Intelligence Planning Systems: Proceedings of the first international conference*. 145–152.

24. Lau, T. L.: 1998, 'Guided Genetic Algorithm'. Ph.D. thesis, Department of Computer Science, University of Essex, UK.
25. Lau, T. L. and E. P. K. Tsang: 1996, 'Applying a mutation-based genetic algorithm to the processor configuration problem'. In: *Proceedings of IEEE 8th International Conference on Tools with Artificial Intelligence*. 17–24.
26. Lau, T. L. and E. P. K. Tsang: 1997, 'Solving the processor configuration problem with a mutation-based genetic algorithm'. *International Journal on Artificial Intelligence Tools* **6**(4), 567–585.
27. Meseguer, P.: 1989, 'Constraint Satisfaction Problems: An Overview'. *AI Communications* **2**(1), 3–17.
28. Mills, P. and E. P. K. Tsang: 2000, 'Guided local search for solving SAT and weighted MAX-SAT problems'. *Journal of Automatic Reasoning, Special Issue on Satisfiability Problems*. Kluwer, **24**, 205–223.
29. Minton, S., M. Johnston, A. Philips, and P. Laird: 1992, 'Minimizing conflicts: A heuristic repair method for constraint satisfaction and scheduling problems'. *Artificial Intelligence* **58**, 161–205.
30. Ruttkay, Z., A. E. Eiben, and P. E. Raue: 1995, 'Improving the Performance of GAs on a GA-hard CSP'. In: *Proceedings, CP95 Workshop on Studying and Solving Really Hard Problems*. 157–171.
31. Selman, B., H. Levesque, and D. Mitchell: 1992, 'A New Method for Solving Hard Satisfiability Problems'. In: *Proceedings of 10th National Conference on Artificial Intelligence*. 440–446.
32. Smith, D. H. and S. Hurley: 1997, 'Bounds for the Frequency Assignment Problem'. *Discrete Mathematics* **167/168**, 571–582.
33. Smith, D. H., S. Hurley, and S. U. Thiel: 1998, 'Improving Heuristics For The Frequency Assignment Problem'. *European Journal of Operational Research* **107**, 76–86.
34. Smith, G. D., A. Kapsalis, V. J. Rayward-Smith, and A. Kolen: 1995, 'Radio link frequency assignment problem report 2.1 - Implementation and testing of genetic algorithm approaches'. Technical report, School of Information Science, University of East Anglia, UK.
35. Tam, V. and P. Stuckey: 1998, 'An Efficient Heuristic-based Evolutionary Algorithm for Solving Constraint Satisfaction Problems'. In: *Proceedings, 3rd IEEE Symposium on Intelligence in Neural and Biological Systems (INBS)*. 21–23.
36. Tiourine, S., C. Hurkens, and J. Lenstra: 1995, 'An overview of algorithmic approaches to frequency assignment problems'. In: *Proceedings of CALMA Symposium, Scheveningen*.
37. Tsang, E. P. K.: 1993, *Foundations of Constraints Satisfaction*. Academic Press Limited.
38. Tsang, E. P. K. and C. Voudouris: 1997, 'Fast local search and guided local search and their application to British Telecom's workforce scheduling problem'. *Operations Research Letters* **20**(3), 119–127.
39. Verfaillie, G., M. Lemaitre, and T. Schiex: 1996, 'Russian Doll Search for Solving Constraint Optimization Problems'. In: *Proceedings, 13th National Conference for Artificial Intelligence*. 181–187.
40. Voudouris, C.: 1997, 'Guided local search'. Ph.D. thesis, Department of Computer Science, University of Essex, UK.
41. Voudouris, C.: 1998, 'Guided Local Search – An Illustrative Example in Function Optimization'. *BT Technology Journal* 46–50.

42. Voudouris, C. and E. P. K. Tsang: 1995, 'Guided local search'. Technical Report CSM-247, Department of Computer Science, University of Essex, UK.
43. Voudouris, C. and E. P. K. Tsang: 1996, 'Partial constraint satisfaction problems and guided local search'. In: *Proceedings of Practical Application of Constraint Technology*. 337–356.
44. Voudouris, C. and E. P. K. Tsang: 1998, 'Guided local search and its application to the travelling salesman problem'. *European Journal of Operations Research* **113**(2), 80–110.
45. Voudouris, C. and E. P. K. Tsang: 1999, 'Solving the Radio Link Frequency Assignment Problem using Guided Local Search'. *Frequency Assignment, Sharing and Conservation Systems (Aerospace)*, Research and Technology Organization (RTO) Meeting Proceedings 13, North Atlantic Treaty Organization (NATO), 14a-1-11
46. Warwick, T.: 1995, 'A GA approach to constraint satisfaction problems'. Ph.D. thesis, Department of Computer Science, University of Essex, UK.
47. Warwick, T. and E. P. K. Tsang: 1994, 'Using a genetic algorithm to tackle the processor configuration problem'. In: *Proceedings of Symposium on Applied Computing*. pp. 217–221.
48. Warwick, T. and E. P. K. Tsang: 1995, 'Tackling car sequencing problems using a generic genetic algorithm'. *Evolutionary Computation* **3**(3), 267–298.