

The COMOGA Method: Constrained Optimisation by Multi-Objective Genetic Algorithms

Patrick D. Surry^{a,b}, Nicholas J. Radcliffe^{a,b}
{pds,njr}@quadstone.co.uk

^aQuadstone Ltd, 16 Chester Street, Edinburgh, EH3 7RA, UK

^bDepartment of Mathematics, University of Edinburgh, King's Buildings, EH9 3JZ, UK

Abstract. This paper describes a novel method for attacking constrained optimisation problems with evolutionary algorithms, and demonstrates its effectiveness over a range of problems. COMOGA (Constrained Optimisation by Multi-Objective Genetic Algorithms) combines two evolutionary techniques for multi-objective optimisation with a simple regulatory mechanism to produce a constrained optimisation method. It shares the universal applicability of penalty-function approaches, but requires significantly fewer free control parameters.

COMOGA takes a dual perspective, considering a constrained optimisation problem sometimes as a constraint satisfaction problem, and sometimes as an unconstrained optimisation problem. These two formulations are treated simultaneously, using a single population, by basing each selection decision on the basis of either constraint violation or function value. A simple adaptive feedback mechanism couples the two formulations by adjusting the relative likelihood of these choices. Unlike penalty function approaches, COMOGA dynamically adapts the emphasis placed on constraint satisfaction and objective function value as the optimisation proceeds, usually yielding final populations which are both feasible and highly fit.

COMOGA has been successfully applied to real industrial problems with comparable performance to highly tuned penalty function approaches. On a test suite of constrained problems previously studied by Michalewicz, application of COMOGA required minimal effort but proved superior to all previous evolutionary methods known to have been applied; indeed it was the only method which found feasible solutions in every run for every problem.

1 Introduction

It is a frequent criticism of evolutionary algorithms that published results are usually obtained with contrived problems without constraints, leading to the suggestion that evolutionary methods are unsuitable for tackling complex constrained optimisation problems. Within the community, there is a wide-spread perception that penalty function methods are a rather blunt instrument for handling general constraints (e.g. Michalewicz, 1992), exhibiting great sensitivity to the values of their many free parameters, and feeding rather too little information back to the algorithm to allow it to handle the constraints satisfactorily. While other methods are available for problems with explicit constraints (including repair methods, Davis & Orvosh, 1993; smart decoders, Davis, 1987, 1991; and special operators incorporating problem knowledge, Michalewicz & Janikow, 1991), these do

not have fully general applicability, and tend to require significant work for each new class of problems tackled. There is thus a need for a method that combines the generality of penalty function approaches with a greater feedback of information to the underlying search algorithm about the way in which progress is being made with the various constraints under consideration.

In this paper we present such a method, COMOGA, based on ideas for multi-objective optimisation. Section 2 presents an introduction to constrained optimisation and constraint satisfaction problems, and surveys the evolutionary techniques which have been used to tackle such problems. Multi-objective optimisation is reviewed in section 3, where the natural fit with population-based algorithms is discussed, and a link is made between multi-criterion optimisation and constraint satisfaction. In section 4 these two strands are drawn together to motivate the COMOGA method, which is then described in detail. The technique is demonstrated for a gas-network problem in section 5, and results for a test suite of constrained optimisation problems previously studied by Michalewicz are summarised in section 6.

2 Constrained Optimisation

2.1 Formulation

Many optimisation problems can be phrased as the minimisation¹ of a given function f , over a search domain \mathcal{S} :

Minimise

$$f : \mathcal{S} \rightarrow \mathbb{R}$$

subject to [the solution $x \in \mathcal{S}$ satisfying certain equalities or inequalities].

The equations or inequalities that the solution x must satisfy are known as constraints. Solutions which satisfy all constraints are said to be *feasible* and the set of all such solutions is called the *feasible region*, \mathcal{S}_F . (Thus the set of optima, \mathcal{S}^* , is a subset of \mathcal{S}_F .) In a *constraint satisfaction problem*, the objective function f is discarded, with the goal being simply to find any solution in \mathcal{S}_F .

Constraints can be characterised in various ways. An inequality constraint is said to be *active* at a point x if it is satisfied as an equality at x . It is typical in constrained optimisation that a number of constraints are active for optimal solutions, so that \mathcal{S}^* is at the boundary of the feasible and infeasible regions (with the result that any weakening of the constraints would change \mathcal{S}^*). Indeed, in highly constrained problems it is often the case that all feasible solutions are near this boundary, and that the volume of the feasible region is negligible when compared with that of the unconstrained search domain. (Although such problems are typically difficult, the converse is clearly not true, with many difficult constrained problems having large feasible regions.)

The constraints on x are conveniently divided into two (imprecise) categories—implicit and explicit. Explicit constraints are those which can be reduced to simple conditions

¹ In this paper we assume, whenever appropriate, without loss of generality, that problems are cast as minimisation problems

on x and are verifiable “by inspection” while implicit constraints are those which specify a condition on some function of x , that requires significant computation to verify (comparable to, or perhaps much greater than, computing $f(x)$). For example, $x_1 < 3$ and $x_1 + x_2 = x_3$ would normally be regarded as explicit constraints while a condition such as ‘the pressure on the wing should not exceed 3N/m^2 ’, where the pressure is computed by a fluid-flow simulation, would be regarded as an implicit constraint. The distinction is useful because genetic operators can usually be constructed that respect explicit constraints whereas this is impractical for implicit constraints.

2.2 Evolutionary Approaches

When a constrained problem is tackled using an evolutionary algorithm, there are several main approaches, with varying degrees of generality (see for example the survey in Michalewicz, 1995b). Perhaps the simplest idea is to restrict the search to the feasible region. This can be done by rejecting infeasible solutions outright, by using greedy decoders or repair mechanisms, or by designed specialised operators which incorporate knowledge of the constraints. The search is thus reformulated as an unconstrained optimisation problem over the reduced space \mathcal{S}_F (the *feasible* region), which is illustrated in figure 1. The diagram shows the image of the search space under the vector-valued function $\mathcal{I} : \mathcal{S} \rightarrow (\mathbb{R}^+)^2 \times \mathbb{R}$ where $\mathcal{I} \triangleq (c_1, c_2, f)$ with $c_i(x)$ measuring the degree to which x violates the i th constraint and $f(x)$ giving its cost (to be minimised).

Rejection of infeasible solutions is generally applicable, but is typically limited in its practical utility, due to the low density of feasible solutions in practical problems (e.g. densities of 1 in 10^{10} are not uncommon). In order to avoid generating and rejecting large numbers of infeasible solutions, *greedy decoders* can be used, in which a problem-specific growth function is designed. Here, the genotype does not directly encode a solution in \mathcal{S} but rather a set of parameters which is used by the decoder to generate a feasible solution. Because the decoder must be guaranteed to never produce infeasible solutions (regardless of the provided parameters), it is often extremely difficult to design. Moreover, it is typically hard to generate a decoder that can be guaranteed to be capable of generating optimal or near optimal solutions.

A related approach is to use repair mechanisms to produce feasible solutions from infeasible ones, mapping $\mathcal{S} \rightarrow \mathcal{S}_F$. (Here, genotypes directly represent solutions in \mathcal{S}). This requires a problem-dependent operator which is able to modify any infeasible solution in such a way as to produce a (nearby) feasible solution. Again this is clearly difficult for many types of constraints. When such mechanisms are employed, a further choice is available, namely whether to write the repaired solution back to the genome, or to use it during evaluation, but then to leave the (infeasible) genome intact. The former approach, which is known as Lamarckism, has the advantage of generally allowing faster local improvement, but can make it harder for the search to traverse infeasible regions of the search space, particularly when \mathcal{S}_F is strongly disconnected with respect to the genetic operators. The latter approach, (which has some parallels with the Baldwin effect) has converse advantages and disadvantages. Davis & Orvosh (1993) present anecdotal empirical evidence that writing the repaired solution back to the genome probabilistically, about 5% of the time, is a good option, and the ideas are further explored by Whitley *et al.* (1994).




Fig. 1. A constrained optimisation problem to minimise cost f while satisfying two implicit constraints is illustrated. Points in the search space \mathcal{S} are shown under the three-dimensional mapping $\mathcal{I} \triangleq (c_1, c_2, f)$. The c_1 and c_2 axes measure the degree of constraint violation, so that points in the feasible region, \mathcal{S}_F , are mapped to the line where both are zero. The desired set of optima \mathcal{S}^* is the set of feasible points with minimum cost. All solutions in \mathcal{S} are mapped to points on and above the shaded surface.

The final suggestion is to build and use genetic operators that “understand” the constraints, in the sense that the syntactic actions of the operators never produce infeasible solutions (ones that violate the constraints). This approach is advocated by Michalewicz (1992), and Radcliffe (1994). Michalewicz & Janikow (1991) have shown how genetic operators can be built that “understand” linear constraints in this sense, and Schoenauer & Michalewicz (1996) construct operators which maintain solutions on nonlinear analytical constraint surfaces; however, it is clear that for many types of constraints (particularly implicit ones) this approach is impractical.

In a problem with implicit constraints, it is often at least as difficult to determine whether a solution is feasible as to evaluate its cost. In such a case it is typically impossible to utilise repair mechanisms or greedy decoders and impractical to restrict search to simply reject infeasible solutions, and we must take an alternative, generally less attractive approach. By exploring infeasible solutions, the goal is to drive the search towards the feasible region. Particularly in problems with many active constraints at optima, we hope to approach optima from ‘both sides’, that is to find nearly feasible solutions with better than optimal function values along with feasible solutions with nearly optimal function values (indeed, this is the explicit goal of the segregated genetic algorithm of LeRiche *et al.*, 1995). Thus it has been stated that:

[g]ood search should approach the optimal from both sides of the feasible/infeasible border (Richardson *et al.*, 1989).

The problem then becomes one of how to compare feasible and infeasible solutions, since in the final analysis it is only feasible solutions which are acceptable. The most widely applicable approach is to employ a penalty function. Here, the search is treated as an unconstrained problem over \mathcal{S} , but the objective function is modified for infeasible solutions by adding terms which degrade their performance. In general, the size of penalty added reflects in some way the *degree* of constraint violation (for example the sum of the constraint violation for each constraint). It is also reasonably standard practice (e.g. Richardson *et al.*, 1989; Michalewicz, 1992) to increase the size of penalties during the course of a run (reverse annealing), so that while a degree of violation is tolerated early in early generations, this tolerance reduces over time. This ensures that, after sufficient time, the optimal solutions to the unconstrained problem using the modified objective function coincide with the optimal solutions to the original constrained problem.

Although penalty functions are essentially universally applicable, they exhibit a number of drawbacks. First, they are weak, in the (formal) sense that they do not provide any problem-specific information to the algorithm. This contrasts with repair mechanisms and problem-specific move operators that exploit understanding of the constraints to provide stronger guidance to the algorithm, but such techniques are not applicable for general constraints. Secondly, the choice of weighting for the constraints is a somewhat subtle matter, particularly when there are many, and increases yet further the number of free parameters to the evolutionary algorithm. Because any choice of parameters determines a fixed form for the modified function in the infeasible region, it induces a fixed ranking on all infeasible solutions. This limits the way in which the search algorithm can explore the infeasible region, since fixed tradeoffs between the various constraints have been introduced. The resulting quality of solution obtained—in fact, the likelihood of finding *any* feasible solution—may be extremely sensitive to the values chosen.

In the next section, we present a method, previously discussed in the context of a specific optimisation problem (Surry *et al.*, 1995), which avoids this difficulty by appealing to the methods of multi-criterion optimisation. Some initial exploration has taken place in this area. For example the work of Schoenauer & Xanthakis (1993) treats each constraint in turn to avoid amalgamating them. Richardson *et al.* (1989) suggest the possibility of using multi-objective techniques (using fitness and either the sum or the number of constraint violations as two objectives) but claim to have been plagued by difficulties. In fact Chu & Beasley (1995) implement a scheme similar to this to deal with a single constraint (using what they term fitness and unfitness). However, they give no guidance in dealing with multiple, non-commensurate constraints, other than by combining them using what is essentially a penalty function (see section 3). The method we propose in section 4 presents a novel solution to this problem.

3 Multi-Objective Optimisation

3.1 Formulation

In many real-world optimisation problems there is not a single objective but a set of criteria against which a solution may be measured. Such problems are often known as

multi-objective or *multi-criterion* optimisation problems, and are defined by a set of objective functions f_1, f_2, \dots, f_N over the search space S , each of which should ideally be minimised.

Perhaps the most common approach to multi-criterion optimisation is to form a new objective function F that is a weighted sum of the individual objectives,

$$F = \sum_{i=1}^N \alpha_i f_i, \quad \alpha_i \in \mathbb{R}^+$$

and to seek to minimise this sum.

If there exists a solution $x^* \in S$ that simultaneously succeeds in minimising each of the f_i , this approach can be reasonably satisfactory, because in this case, successful optimisation of F will also optimise each f_i . In the more general case, however, the component objectives f_i will *compete*, in the sense that improvement against one will in some cases require a degradation against another. In this case, the approach of forming a weighted sum is less attractive, because the choice of weights α_i will determine the trade-off between the various component objectives that optima of the combined function F will exhibit. This is particularly unsatisfactory in cases where the various objectives are *non-commensurate*, in the sense that trade-offs between them are either arbitrary or meaningless. A good example of this might arise when seeking to maximise profit while minimising ecological damage, where most people would accept that any assignation of economic cost to ecological damage is at best arbitrary.

In the case of multi-objective problems with competing, non-commensurate criteria, a more satisfactory approach is to search not for a single solution but for that set of solutions that represent the “best possible trade-offs.” Such solutions are said to be *Pareto-optimal*, (after Vilfredo Pareto who first advanced the concept) and are characterised by introducing the notion of domination. A solution x is said to *dominate* another solution y if its performance against each of the objective functions is at least as good as that of y , and its performance is better against at least one objective, i.e. if and only if

$$\begin{aligned} \forall i \in \{1, 2, \dots, n\} : & \quad f_i(x) \leq f_i(y) \\ \text{and } \exists j \in \{1, 2, \dots, n\} : & \quad f_j(x) < f_j(y). \end{aligned}$$

Clearly in this case, x may reasonably be said to be a superior solution to y . The *Pareto-optimal set* (or *front*) \mathcal{P} is the set of solutions that are not dominated by any other solution in the search space, i.e.

$$\mathcal{P} \triangleq \{x \in S \mid \nexists y \in S : y \text{ dominates } x\}.$$

3.2 Evolutionary Approaches

Although it is difficult to attack multi-criterion problems with traditional optimisation methods, it is relatively natural in population-based search algorithms to consider trying to use the population to hold solutions that represent different trade-offs. Reasonably simple modifications to the selection (and perhaps the replacement) method may be all that is required to effect this. A number of schemes have been proposed, most

of which are based around the notion of only allowing selective advantage between solutions when one dominates another. Fonseca & Fleming (1995) provide an overview of many such techniques. The effectiveness of these methods is further enhanced when combined with some form of niching, to encourage greater diversity in the population. Niching methods include structured population models (e.g. Norman, 1988; Manderick & Spiessens, 1989; Gorges-Schleuter, 1989), sharing (Goldberg & Richardson, 1987), and crowding (Cavicchio, 1970; DeJong, 1975).

3.3 Constraint Satisfaction as Multi-Criterion Optimisation

It is clear that the constraint satisfaction problem is (formally) equivalent to the simple class of multi-objective problems (discussed above) in which all objectives can be minimised simultaneously. We measure the degree of constraint violation for each constraint (or group of commensurate constraints) and treat each of these as an objective in a multi-criterion problem. (Consider again figure 1, where we are now only required to find a solution in \mathcal{S}_F . The dashed lines, if extended upwards as vertical manifolds, might indicate a series of progressively dominating surfaces, converging on $\mathcal{I}(\mathcal{S}_F)$ —the Pareto-optimal set in this case.) Although, in such a case, minimising a penalty function expressing the degree of constraint violation would be the most common approach, we suggest that a more appropriate strategy when using evolutionary techniques is to use the simple techniques for general multi-criterion optimisation discussed above in order to exploit the ability of the population to hold many different possible trade-offs between the constraints. This allows the algorithm to dynamically “discover” an appropriate trajectory by which to approach the feasible region, rather than arbitrarily assigning the relative importance of different combinations of constraint violations.

In the next section we show how this idea can then be extended to the constrained optimisation problem, where not only must several constraints be satisfied, but a given objective function f must also be minimised.

4 The COMOGA Approach

4.1 Motivation

It was pointed out in section 2 that if we choose to use a penalty function with some given set of parameters to attack a constrained optimisation problem, we make an *a priori* decision about the relative importance of different degrees of constraint violation, regardless of their actual difficulty to satisfy. Further, by combining the degree of constraint violation with the objective function value we impose fixed choice concerning the tradeoff between constraint satisfaction and optimisation.

In section 3, we showed that the methods of evolutionary multi-objective optimisation can be applied directly to the constraint satisfaction problem, and proposed that we could apply similar ideas to constrained optimisation. Obviously, the situation is complicated somewhat by the additional requirement of minimising some function over the feasible region. Here we think of f as an extra criterion which is of less importance than any of the “constraint criteria”, i.e. there is no acceptable trade-off between minimising (satisfying) the constraints, and minimising f .

The approach we advocate is to view a constrained optimisation problem alternatively as a constraint satisfaction problem (ignoring the objective function) and as an unconstrained optimisation problem (ignoring the constraints). We then propose to decide adaptively which view to take at any instant based on the relative success with respect to the two formulations. In order to find near-optimal solutions, we must be careful to get neither “too far” from feasibility nor “too far” from optimal fitness, while also recognising that constraint satisfaction is more important than optimisation (as ultimately we are only interested in feasible solutions). We show that an adaptive population-based algorithm is ideal for this purpose.

There are numerous approaches to unconstrained optimisation using a population-based algorithm, and various techniques for constraint-satisfaction based on multi-criterion evolutionary algorithms have been discussed above. The difference between these two types of algorithms can typically be ascribed solely to the selection (and replacement) regime—in the first case selective decisions are based on fitness (cost) while in the second they are normally based on some form of Pareto ranking.

This motivates an attractively simple scheme for coupling the combined constrained optimisation problem, in which we use a single algorithm but randomly decide each time a selective decision must be made whether to consider the problem as a constraint satisfaction problem or as an unconstrained optimisation problem. We then adjust the relative likelihood of adopting each view using a simple feedback mechanism that tries to maintain a fixed fraction of the population in the feasible region. Because individual solutions can be selected on the basis of either constraint satisfaction or cost, this results in an algorithm which aims to dynamically explore the boundary between feasibility and infeasibility *without* arbitrary penalty factors fixing the relative quality of the different achievable tradeoffs.

This can be seen as a generalisation of the scheme recently proposed by Chu & Beasley (1995) which can importantly handle more than one constraint without having to amalgamate them.

4.2 Algorithm

To treat the constraint satisfaction aspect of the problem, we can conceptually label all members of the search space \mathcal{S} with some measure of their Pareto ranking based on constraint violation, either by conceptually peeling off successive non-dominating layers (Goldberg, 1989), or by assigning to each solution a “rank” equal to the number of solutions which dominate it (Fonseca & Fleming, 1993). (The latter scheme has the advantage that it is easy to subtract the effect of a deleted individual and add the effect of a new individual without re-ranking the entire population.) Note that this ranking is a dynamic one, based on the current population of achievable constraint tradeoffs rather than a fixed ranking of any possible combination of constraint violations. We denote this *population-dependent* ranking function $R : (\mathbb{R}^+)^N \rightarrow \mathbb{Z}^+$, where N is the number of constraints.

From the unconstrained optimisation view, every solution has some cost value associated with it. Thus we are presented with a dual view of each solution in the population and can form the two-dimensional mapping $\mathcal{I}_R : \mathcal{S} \rightarrow \mathbb{Z}^+ \times \mathbb{R}$, with $\mathcal{I}_R \hat{=} (R \circ (c_1, \dots, c_N), f)$. This reduces the problem to the two-objective problem illustrated in

figure 2. We must couple the two viewpoints in order to solve the combined constrained optimisation problem: in solving the constraint satisfaction problem we minimize along the R axis and in solving the unconstrained optimisation problem we minimize along the f axis. However, we desire not simply solutions on the Pareto-optimal surface \mathcal{P} , but rather solutions in the intersection of the Pareto-optimal set with the feasible region (as constraint satisfaction is “more important” than cost minimisation).




Fig. 2. The constrained optimisation problem with N constraints can be re-cast as a two-objective problem by assigning a Pareto rank based on constraint violation. The Pareto-optimal set \mathcal{P} is the set of non-dominated solutions under $\mathcal{I}_R \triangleq (R \circ (c_1, \dots, c_N), f)$. The feasible set is mapped to the line segment $\mathcal{I}_R(\mathcal{S}_F)$, and the desired set of optima is mapped to their intersection, $\mathcal{I}_R(\mathcal{S}^*)$. The search space \mathcal{S} is mapped to points on and above \mathcal{P} .

One possible approach is to use a sub-ranking scheme, where only solutions with equal Pareto rank for constraints are distinguished on the basis of cost. However, this is likely to result in an evolutionary process which first concentrates on the constraint satisfaction problem (hence sampling points in the feasible region essentially at random) and only once this is solved tries to reduce cost. This “approach from above” not only lacks the desirable property of being able to combine low-cost, nearly-feasible solutions with higher-cost feasible ones, but may be an extremely poor way to search \mathcal{S}_F if it is a highly sparse and disconnected subset of \mathcal{S} .

An appealing alternative approach is to enlist the ideas of Schaffer (1985). In his *vector evaluated genetic algorithm* (VEGA), he selects some fraction (typically $1/k$) of the population based on each of the k objective functions. When a fixed fraction is used for each objective (e.g. $1/k$), this tends to favour the development of “specialist” populations that excel in one objective function, particularly when fitness-proportionate

selection is used (Richardson *et al.*, 1989). However, COMOGA will actively exploit this tendency by adaptively changing the likelihood of selecting with respect to each objective.

The suggestion in our case is to use, for example, tournament selection (Goldberg, 1989), sometimes basing the tournament on cost f and sometimes on the Pareto ranking R with respect to constraint violation. (In cases where the selected attributes are equal, the other attribute is compared.) A probability p_{cost} is used to determine the likelihood of selection with respect to cost, and will be adapted as the algorithm progresses. Any fixed value of p_{cost} will induce an overall probability of reproduction equal to some linear combination of the reproductive probabilities with respect to the two attributes, with *population-dependent* weights. Although such a fixed p_{cost} may favour convergence to some non-feasible point on the Pareto-optimal curve, it is clear that as $p_{cost} \rightarrow 0$, the process increasingly favours constraint rank until in the limit of $p_{cost} = 0$ we are essentially solving the constraint-satisfaction problem; seeking feasible solutions regardless of cost (unless the constraint rankings are equal—this is equivalent to the sub-ranking approach described above). We thus hope that some intermediate non-zero value will allow us to find feasible solutions of low cost. This is illustrated in figure 3.




Fig. 3. Using a VEGA-like scheme of selecting probabilistically with respect to one of the two objectives (cost or constraint rank), we induce a perceived fitness of some population-dependent weighted combination of the two objectives. As p_{cost} tends to zero, the scheme favours constraint rank more, and cost less. By adaptively changing p_{cost} based on the proportion of feasible solutions observed in the population, the algorithm dynamically discovers how to achieve constraint satisfaction and minimisation simultaneously.

To avoid the problem of fixing a particular value for p_{cost} , we propose to change the value adaptively by setting a target proportion τ of feasible solutions in the population. (τ is similar to the flip threshold of Schoenauer & Xanthakis, 1993.) We start by choosing some arbitrary value for p_{cost} , say 0.5, and some desired proportion of feasible solutions, e.g. $\tau = 0.1$. After each generation, if the number of feasible solutions in the population is not close to τ , then we adjust p_{cost} up or down accordingly: if the actual proportion is too low, we decrease p_{cost} , e.g.

$$p_{cost} \leftarrow (1 - \varepsilon)p_{cost}, \quad (1)$$

and conversely, if the proportion is too high, we increase it, e.g.

$$p_{cost} \leftarrow 1 - (1 - p_{cost})(1 - \varepsilon). \quad (2)$$

This does, of course, introduce several new parameters to the algorithm (though notably fewer than a penalty function), which we were trying to avoid, but we find in practice that the scheme is remarkably robust to them, in contrast to the sensitivity of penalty function parameters. This leads to the COMOGA method, which is summarised below.

The COMOGA Method

1. Calculate constraint violations (c_1, \dots, c_N) for all solutions.
2. Pareto rank based on constraint violations (e.g. by counting the number of members of the population dominated by each solution).
3. Evaluate the cost (fitness) of solutions.
4. Select an (expected) proportion p_{cost} of parents based on cost, and the others based on constraint ranking.
5. Apply the genetic move operators (recombination, mutation etc.)
6. Replace an (expected) proportion p_{cost} of solutions based on cost, and the others based on constraint ranking.
7. Adjust p_{cost} if the proportion of feasible solutions in the population is not close to the target proportion, τ , according to equations (1) and (2). Lowering p_{cost} favours feasible solutions and raising it favours lower cost solutions.

Typical values for the parameters are $\tau = 0.1$ and $\varepsilon = 0.1$, with $p_{cost} = 0.5$ initially.

4.3 Discussion

The COMOGA scheme has several attractive features. First, and foremost, it removes the necessity for the many parameters of a penalty function which must be determined empirically. Secondly, it turns the acknowledged weaknesses of VEGA to favour extreme solutions to advantage, as in this case we are only ultimately interested in solutions which excel at constraint satisfaction (have low constraint rank). Thirdly, the adaptive approach to specifying p_{cost} allows the algorithm to find its own trajectory to approach the desired optimal values.

In practice, it is important to incorporate explicit constraints (e.g. linear or other specific nonlinear ones) where possible, and to amalgamate multiple commensurate constraints in order to reduce the dimensionality of the Pareto-optimal front. It may also

be advantageous to incorporate niching or other diversity promoting measures in the algorithm (in the work presented here we have simply enforced uniqueness which is not likely to be highly effective for real parameter optimisation).

5 An Illustrative Application

We will illustrate the application of the COMOGA approach to a gas-network pipe-sizing problem, contrasting the results with a penalty-function approach. This work has been previously reported in greater detail (Surry *et al.*, 1995).

The problem involves determining the diameters of pipes in a fixed-topology network (with fixed supplies and demands), in order to minimise expenditure, while satisfying two implicit constraints defining the minimal pressure at each node along with an engineering requirement that each pipe should have at least one upstream pipe of the same or greater diameter.

In the particular problem considered, the network contained 25 pipes, each of which could be selected from six possible sizes, giving rise to a search space of size $6^{25} \simeq 3 \times 10^{19}$. The network is a real one, which was actually built, with pipe sizes determined using a greedy heuristic method. (Both evolutionary methods outperformed the existing greedy heuristic for the problem, giving identical results.) The density of valid networks (ones which satisfy the constraints) in the search space is extremely low—random sampling of more than 3×10^7 points produced only a single admissible configuration.

Because of the implicit nature of the constraints in the pipe-sizing problem, the only applicable conventional approach is to use a penalty function, as it would be extremely difficult to construct genetic operators that respected them, and prohibitively expensive to use a repair mechanism (if indeed one could be constructed). In order to compare this approach with the COMOGA technique of modifying the selection regime, a standard steady-state elitist duplicate-free evolutionary algorithm using an integer-valued representation and standard recombination and mutation operators was employed (Surry *et al.*, 1995). Binary tournament selection with parameter 1.0 was used to select parents, and the resulting child was re-inserted using a replace-worst scheme. Tournament replacement was also investigated, but the more aggressive replace-worst strategy proved superior.

In the first case, an annealed penalty function which combined a time-dependent weighted sum of the degree of constraint violation for the two constraints along with the basic cost function value was used as the objective. This involved six control parameters to incorporate the two constraints. A wide range of penalty function parameters were tested to discover the typical relative values of constraint violations, etc. As has been widely reported previously, the quality of the resulting algorithms is highly sensitive to these values, with small changes often resulting in runs in which no feasible solution was found.

The technique (with good parameters) produced consistently good results, although it did not always converge to the same optimal solution. In most cases it found networks which were better, often significantly so, than that determined by the heuristic approach. In almost all cases the algorithm found a valid network by the end of the run (i.e. one

in which the penalty terms were zero). A snapshot of a single successful run using the penalty function is discussed in figure 4.

The same algorithm was then adapted to COMOGA approach, as described in section 4. Each member of the initial population was assigned a rank according to constraint violation by counting the number of members in the population by which it was dominated. The same selection and replacement regime was used, but with decisions based on cost value with probability p_{cost} , and otherwise on constraint ranking.

As with the penalty function approach, a variety of population sizes, mutation and crossover rates, and so forth were investigated. Results were best with populations of about 100 individuals, and with the same stopping conditions, runs lasted for similar numbers of evaluations, and produced similar quality solutions (the same “best” solution from the penalty-function approach was found consistently). In contrast to the penalty-function approach however, algorithm performance was much less sensitive to these control parameters. Most importantly, the COMOGA scheme was not particularly sensitive to the method used for adapting the p_{cost} parameter, nor to the target proportion of feasible solutions, τ . A sample run of the COMOGA algorithm is discussed in figure 5.

Although the overall performance of the COMOGA algorithm was very similar to that of the best penalty function approach found, both in terms of computational effort required and frequency of finding the best solutions. However, significantly less experimentation was required to find values for COMOGA’s parameters that work well than was the case with the penalty function method.

6 Experimental Results

In order to validate the COMOGA method it has been applied to a series of known test problems in constrained optimisation for which other evolutionary methods have been applied. The first such problem is the so-called “bump problem” of Keane (1996b), a highly multi-modal maximisation problem defined for an arbitrary number of variables, n , and two non-linear constraints, with an unknown optimum value. For comparative purposes, we studied the problem with $n = 50$. Using an untuned implementation of the COMOGA method, we achieved results superior to any generic evolutionary algorithm previously studied, with fitnesses in the range 0.814 to 0.828 after 200 000 evaluations. Keane (1996a) reports that the best result known to him is 0.832, produced by a non-genetic technique, but more recently Michalewicz & Schoenauer (1996) have achieved 0.833 with a population of size of 30 over 30 000 generations using a problem-specific crossover operator to search the constraint surface, and in this volume Wodrich & Bilchev (1997) report results of 0.826 after only 30 000 evaluations, also employing a problem-specific heuristic.

Michalewicz (1995a) has proposed a test set consisting of five constrained optimisation problems, for which he has compared six existing evolutionary techniques. This test set was coded as a library for the *Reproductive Plan Language*, RPL2 (Surry & Radcliffe, 1994, Radcliffe & Surry, 1994), and an untuned implementation of COMOGA was applied to the problems. Table 1 summarises the results of these experiments in the same form presented by Michalewicz (based on the same number of total evaluations per

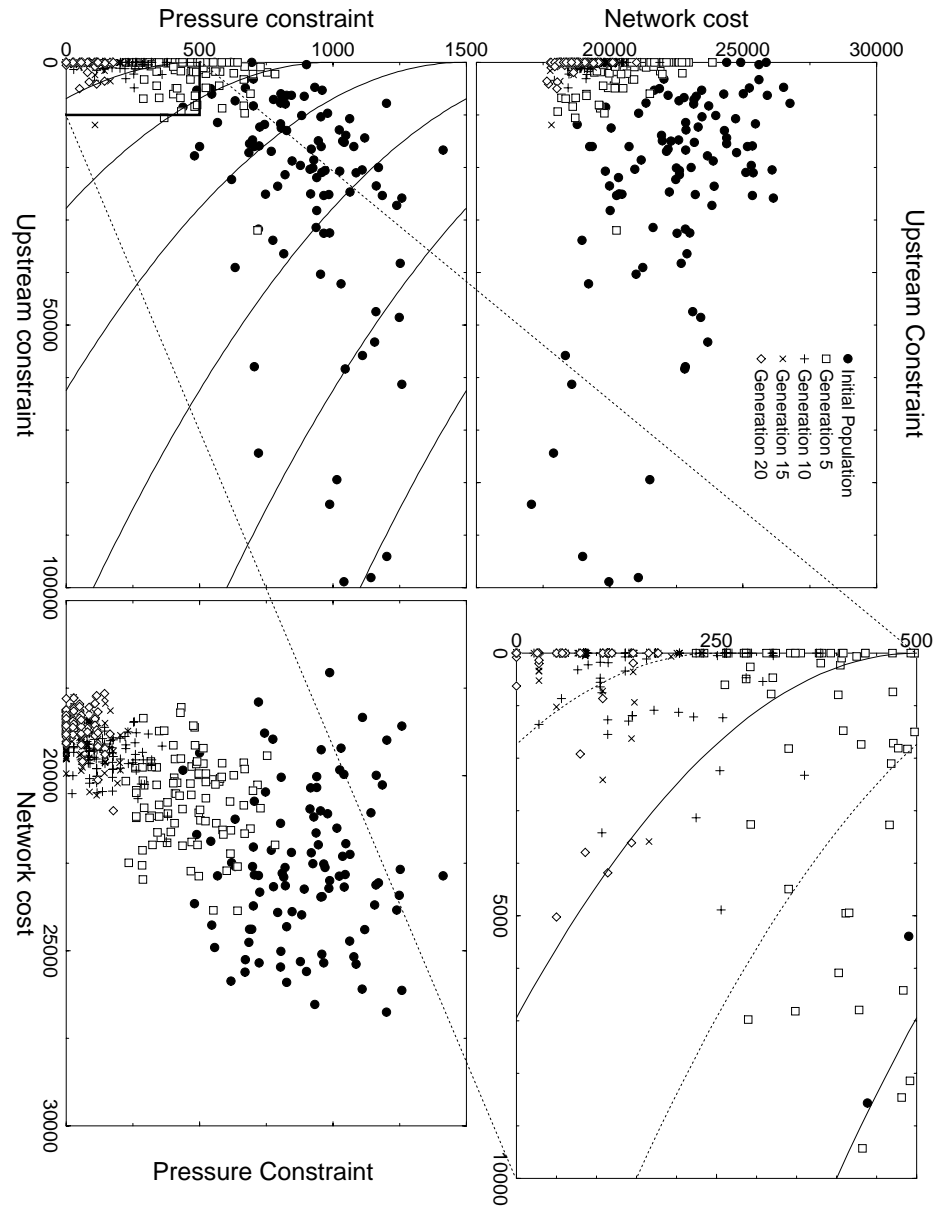


Fig. 4. The figure shows the pattern of convergence for a single run of an algorithm based on a penalty function. A snapshot of the population is shown every five generations, plotted with respect to the degree of violation of the two constraints and to unpenalised cost (to be minimised). An enlarged section of the region near zero constraint violation is also shown, along with curves which show lines of equal penalty when constraint violations are incorporated with the cost. Note that the population approaches the minimum by first satisfying the constraints and then minimising in the feasible region.

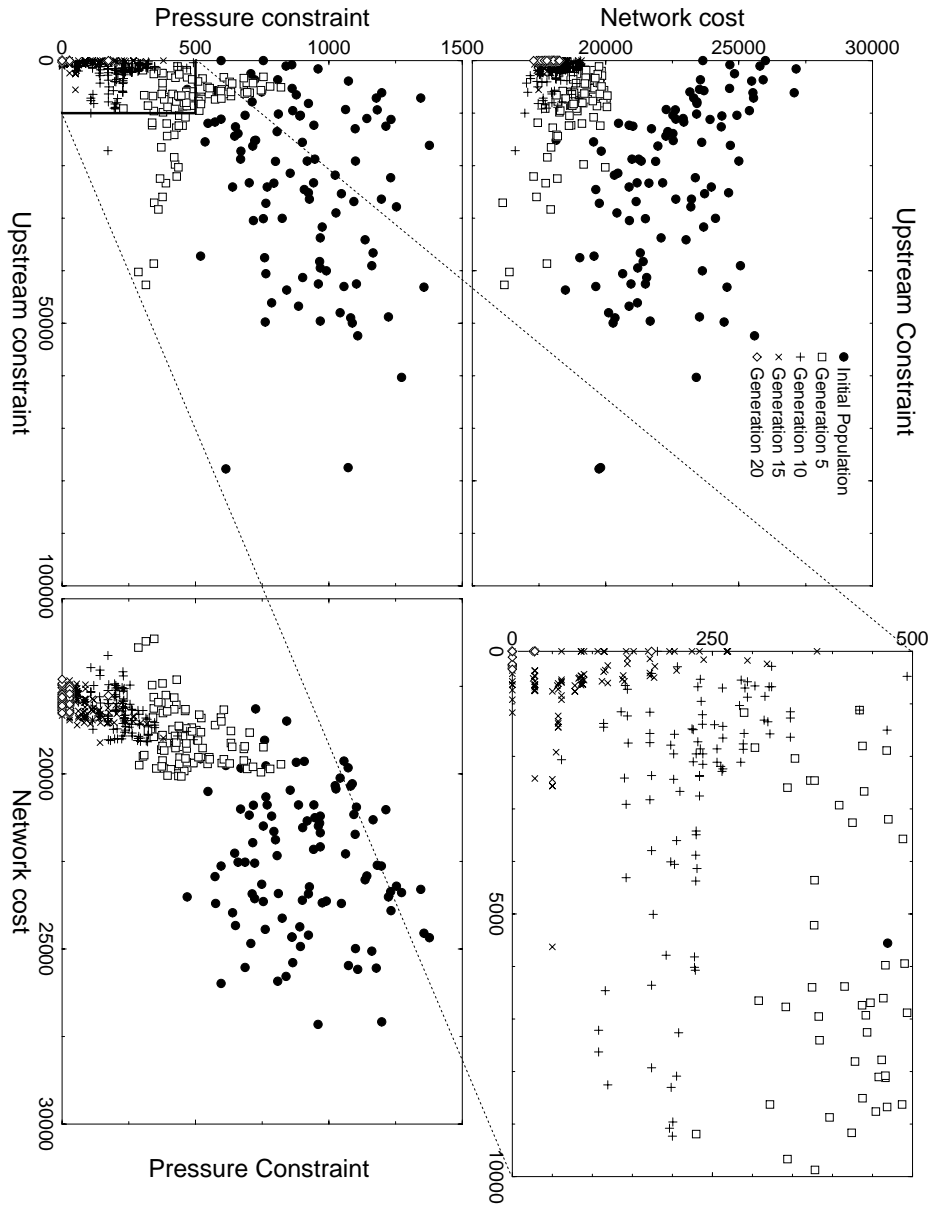


Fig. 5. The figure shows the pattern of convergence for a single run of the COMOGA method. The population is shown every five generations with respect to the two constraints and the cost. In contrast to the penalty function run, the population explores a variety of tradeoffs between constraint satisfaction and cost, approaches the feasible solutions with minimal cost from both above and below.

run). Here all runs resulted in feasible solutions so Michalewicz's summary by degree of violation has been dropped.

Case	Form	Vars	LI	NI	Eq	Optimum	Best/Median/Worst		Rank	
							Previous	COMOGA		
1	quadratic	13	9	-	-	-15.000	#4	-15.000 -15.000 -15.000	-14.997 -14.996 -14.994	5=
2	linear	8	3	3	-	7049.331	#4	7377.976 8206.151 9652.901	7081.43 7556.85 8322.51	1
3	polynomial	7	-	4	-	680.630	#4	680.642 680.718 680.955	680.663 680.690 680.755	1=
4	nonlinear	5	-	-	3	0.054	#4	0.054 0.064 0.557	0.058 0.205 0.570	3
5	quadratic	10	3	5	-	24.306	#2	25.486 26.905 42.358	24.340 24.509 24.710	1

Table 1. The table shows the results of applying the COMOGA method to the test set of Michalewicz (1995a), in which he compared six evolutionary constraint handling methods. The form, number of variables and number and types of constraints (linear inequalities, LI, non-linear inequalities, NI, and non-linear equalities, Eq) are shown for each test problem, along with the optimal value. Results from a set of ten runs are shown for both the previous best method and the COMOGA method, along with the overall ranking on minimum, median and maximum for this new scheme (where this is ambiguous a tie has been awarded). COMOGA performs well on most of the problems, and was the only method to produce a feasible solution in every run for every problem. It performs worst with the linear constraints of problem 1, though this is essentially only a failure to sufficiently polish the parameter values. In problem 4 the equality constraints are treated by essentially converting them to inequality constraints specifying an allowable violation of the equality (a tolerance of 0.001 was used, as this defined feasibility in Michalewicz's study). This creates an artificial division between the feasible and infeasible region.

For problem #1, it has failed only to accomplish the final 'polishing' having achieved 4+ figure accuracy in all of the parameter values for every run. This is more suggestive of shortcomings in the genetic operators used (e.g. perhaps the mutation rate or loss of diversity) rather than a failure in the COMOGA technique itself. Note also that no operators which 'understood' these linear constraints were used for the purposes of this initial comparison.

For problem #4, by treating the equality constraints as hard inequality constraints (using the arbitrary 0.001 factor for allowable degree of violation) we create a sharp artificial boundary between the feasible and infeasible regions. Since this region is essentially a two-dimensional subspace of the five-parameter domain it is perhaps not surprising that it is difficult to “explore” it effectively. In fact, we typically observe that the trajectory of the initial population to find the first feasible solution has a large impact on the quality of the best feasible solution found. This suggests that perhaps this is not the best way in which to incorporate equality constraints within the COMOGA framework, but this requires further investigation.

7 Summary

A new approach of general applicability to constrained optimisation—the COMOGA method—has been presented. This technique treats a constrained optimisation problem sometimes as a constraint satisfaction problem and sometimes an unconstrained optimisation problem, using a single population, by switching between selection regimes. A simple feedback mechanism is used to determine the (expected) proportion of selective decisions which are made with respect to the two viewpoints, depending on the relative progress observed on each.

The COMOGA method uses the memory implicit in the population to “discover for itself” the relative utility of different achievable combinations of constraints and objective function value. The population thus forms not just a pool of good solutions among which recombination takes place, but a context in which to determine the fitness of any one member—the effective weighting of the various constraints is determined by the population, as is the relative weighting of constraint satisfaction and cost minimisation. This contrasts with a penalty function approach, where both are determined *a priori*, and appears to carry the significant benefit of reducing both the sensitivity of the genetic algorithm to the values of the free parameters, and the number of those parameters.

A series of experiments have validated the COMOGA approach, based on the simple yet powerful idea that by merely alternating between two selection (and replacement) regimes we can couple solution of the constraint satisfaction problem to the simultaneous solution of the unconstrained optimisation problem, and discover good solutions to the constrained optimisation problem. Because the scheme is based only on modifying the selection regime, it is possible to use whatever representation, genetic operators and update strategy is appropriate for a given problem within the COMOGA framework.

References

- D. J. Cavicchio, 1970. *Adaptive Search Using Simulated Evolution*. PhD thesis, University of Michigan.
- P. C. Chu and J. E. Beasley, 1995. A genetic algorithm for the set partitioning problem. Technical report, The Management School, Imperial College. (Revised 1997).
- L. Davis and D. Orvosh, 1993. Shall we repair? Genetic algorithms, combinatorial optimization, and feasibility constraints. In S. Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, page 50. Morgan Kaufmann (San Mateo).
- L. Davis, 1987. *Genetic Algorithms and Simulated Annealing*. Pitman (London).
- L. Davis, 1991. *Handbook of Genetic Algorithms*. Van Nostrand Reinhold (New York).
- K. A. De Jong, 1975. *An Analysis of the Behaviour of a Class of Genetic Adaptive Systems*. PhD thesis, University of Michigan.
- C. M. Fonseca and P. J. Fleming, 1993. Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. In S. Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 416–423. Morgan Kaufmann (San Mateo).
- C. M. Fonseca and P. J. Fleming, 1995. An overview of evolutionary algorithms in multiobjective optimization. *Evolutionary Computation*, 3(1):1–16.
- D. E. Goldberg and J. Richardson, 1987. Genetic algorithms for multimodal function optimisation. In *Proceedings of the Second International Conference on Genetic Algorithms*, pages 41–49. Lawrence Erlbaum Associates (Hillsdale).
- D. E. Goldberg, 1989. *Genetic Algorithms in Search, Optimization & Machine Learning*. Addison-Wesley (Reading, Mass).
- M. Gorges-Schleuter, 1989. ASPARAGOS: an asynchronous parallel genetic optimization strategy. In *Proceedings of the Third International Conference on Genetic Algorithms*, pages 422–427. Morgan Kaufmann (San Mateo).
- A. J. Keane, 1996a. (personal communication).
- A. J. Keane, 1996b. *Modern Heuristic Search Methods*, chapter A Brief Comparison of Some Evolutionary Optimization Methods, pages 255–272. J. Wiley.
- R. G. LeRiche, C. Knopf-Lenoir, and R. T. Haftka, 1995. A segregated genetic algorithm for constrained structural optimization. In L. J. Eshelman, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 558–565. Morgan Kaufmann (San Francisco).
- B. Manderick and P. Spiessens, 1989. Fine-grained parallel genetic algorithms. In J. D. Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms*, pages 428–433, San Mateo. Morgan Kaufmann Publishers.
- Z. Michalewicz and C. Z. Janikow, 1991. Handling constraints in genetic algorithms. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 151–157. Morgan Kaufmann (San Mateo).
- Z. Michalewicz and M. Schoenauer, 1996. Evolutionary algorithms for constrained parameter optimization problems. *Evolutionary Computation*, 4(1):1–32.
- Z. Michalewicz, 1992. *Genetic Algorithms + Data Structures = Evolution Programs*. Springer Verlag (Berlin).
- Z. Michalewicz, 1995a. Genetic algorithms, numerical optimization, and constraints. In L. J. Eshelman, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 151–158. Morgan Kaufmann (San Francisco).
- Z. Michalewicz, 1995b. A survey of constraint handling techniques in evolutionary computation methods. In *Proceedings of the 4th Annual Conference on Evolutionary Programming*, pages 135–155. MIT Press (Cambridge, MA).
- M. Norman, 1988. A genetic approach to topology optimisation for multiprocessor architectures. Technical report, University of Edinburgh.

- N. J. Radcliffe and P. D. Surry, 1994. The Reproductive Plan Language RPL2: Motivation, architecture and applications. In J. Stender, E. Hillebrand, and J. Kingdon, editors, *Genetic Algorithms in Optimisation, Simulation and Modelling*, pages 65–94. IOS Press (Amsterdam).
- N. J. Radcliffe, 1994. The algebra of genetic algorithms. *Annals of Maths and Artificial Intelligence*, 10:339–384.
- J. T. Richardson, M. R. Palmer, G. Liepins, and M. Hilliard, 1989. Some guidelines for genetic algorithms with penalty functions. In *Proceedings of the Third International Conference on Genetic Algorithms*, pages 191–195. Morgan Kaufmann (San Mateo).
- J. D. Schaffer, 1985. Multiple objective optimization with vector evaluated genetic algorithms. In *Proceedings of an International Conference on Genetic Algorithms*, pages 93–100. Lawrence Erlbaum.
- M. Schoenauer and Z. Michalewicz, 1996. Evolutionary computation at the edge of feasibility. In H.-M. Voigt, W. Ebeling, I. Rechenberg, and H. Schwefel, editors, *Parallel Problem Solving from Nature IV*, pages 245–254. (Springer-Verlag, LNCS 1141).
- M. Schoenauer and S. Xanthakis, 1993. Constrained GA optimisation. In S. Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 573–580. Morgan Kaufmann (San Mateo, CA).
- P. D. Surry and N. J. Radcliffe, 1994. RPL2: A language and parallel framework for evolutionary computing. In Y. Davidor, H.-P. Schwefel, and R. Männer, editors, *Parallel Problem Solving from Nature III*, pages 628–637. Springer-Verlag, Lecture Notes in Computer Science 866.
- P. D. Surry, N. J. Radcliffe, and I. D. Boyd, 1995. A multi-objective approach to constrained optimisation of gas supply networks: The COMOGA method. In T. C. Fogarty, editor, *Evolutionary Computing: AISB Workshop*, pages 166–180. Springer-Verlag, Lecture Notes in Computer Science 993.
- D. Whitley, V. S. Gordon, and K. Mathias, 1994. Lamarckian evolution, the baldwin effect and function optimization. In Y. Davidor, H.-P. Schwefel, and R. Männer, editors, *Parallel Problem Solving from Nature III*, pages 6–15.
- M. Wodrich and G. Bilchev, 1997. Cooperative distributed search: the ants' way. (this volume).