Chapter 1
Understanding the impact of constraints: a rank based fitness function for evolutionary methods

Eric S. Fraga and Oluwamayowa Amusat

1.1 Introduction

Model based process design is often formulated as an optimisation problem. The problem definition includes one or more objective functions and both equality and inequality constraints. In process design, many of the constraints originate from underlying physical laws. For instance, the temperature of a distillation plate at equilibrium is described by Raoult's Law, an equality constraint, or the amount of mass in a vessel must be greater than 0, an inequality constraint. These constraints cannot be violated if the design obtained is to be realisable physically.

However, there are other types of constraints. Some constraints are of the form of

It would be great if the solution we obtained had this characteristic.

Examples include: the temperature of the room in the dwelling should be 20 °C; the purity of this by-product should be at least 90%; the pressure changes should be less than some amount specified; and so on. Constraints such as these can be reformulated as objectives and then either incorporated into a single objective function using a penalty term or the problem is transformed into a multi-objective problem. The use of penalty terms (or weighted objective functions equivalently for multi-objective formulations) is difficult due to the need to choose the penalty weights. Defining the problem as a multi-objective optimisation problem and using multi-objective solution methods is therefore more attractive.

There are many methods for multi-objective optimisation; see, for instance, Coello Coello and Landa Becerra (2009) for a selected review, concentrating on

Eric S. Fraga
Centre for Process Systems Engineering, Department of Chemical Engineering, UCL (University College London), e-mail: `e.fraga@ucl.ac.uk`

Oluwamayowa Amusat
Centre for Process Systems Engineering, Department of Chemical Engineering, UCL (University College London), e-mail: `oluwamayowa.amusat.13@ucl.ac.uk`

evolutionary methods. Multi-objective optimisation methods will traditionally attempt to generate a population of solutions that consist of *non-dominated* solutions and therefore represent an approximation to the Pareto front (Pardalos et al, 2016), assuming that they are able to identify globally optimal solutions (Törn and Žilinskas, 1989). As such, in principle, any multi-objective method is suitable for tackling the types of problems noted above. On the other hand, multi-objective problems that are derived from relaxing a *desirable* constraint are subtly different from more general multi-objective problems: the original single objective is somehow more important. Therefore, it is desirable to have a multi-objective optimisation method that emphasises solutions that have more favourable values for that objective. The aim is to provide the design engineer with insight into how the relaxation of the desirable constraint affects the main objective function.

The aim of this chapter is to present a fitness function that provides the selection pressure on population based evolutionary optimisation methods to generate solutions with a preference for those that improve the main objective function but not at the expense of ignoring the other objectives completely. We illustrate this fitness function using a new multi-objective evolutionary method adapted from an existing single objective *plant propagation algorithm* (Salhi and Fraga, 2011).

The novel multi-objective optimisation method is presented and is applied to the problem of designing an integrated energy generation and storage system for off-grid mining operations. Although the design objective is to minimise capital cost, there is a requirement imposed to design the system so that it needs no fuel brought in from off-site. This requirement is a desired property and in this chapter is treated as a separate objective. The results presented show that a multi-objective approach does provide for a better understanding of what is possible and hence what may be desirable as the final solution.

1.2 A Multi-objective Rank Based Fitness Function for Pareto Extremes

Consider the set of points shown in Fig. 1.1 which plots the points in the space of objective function values for a bi-criteria problem. The assignment of fitness values to these points can be done in a variety of ways. Typically, the non-dominated solutions will all be given the same fitness value, the best fitness when compared with the fitness assigned to dominated points. The dominated points will then be assigned fitness values in different ways. One approach is to remove the non-dominated points from the set, find the now non-dominated points and assign these all the same fitness value, one that is worse than that assigned to the original non-dominated points. Then repeat the same process until no points are left. An alternative is to assign a fitness to the originally dominated points based on the distance of these points to either an approximation to the Pareto front defined by the piecewise linear fit to the non-dominated points or by the distance to the nearest non-dominated point (Fiandaca et al, 2009). The aim of these fitness methods is to emphasise the non-dominated points and hence drive an evolutionary algorithm towards a good approximation to the Pareto front.

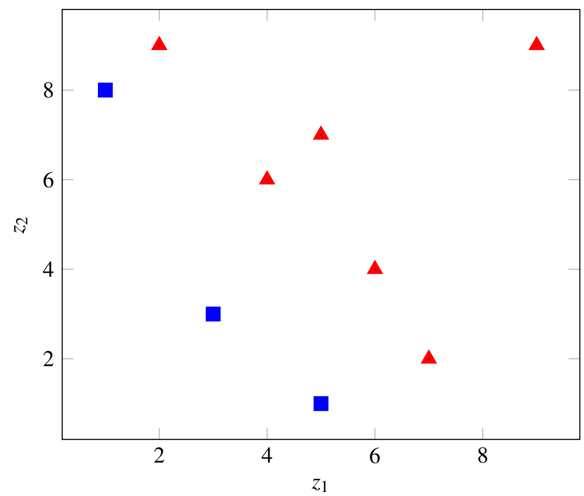

Fig. 1.1 Simple scatter plot with non-dominated points indicated by blue squares and dominated points by red triangles, assuming that the goal is to minimise both objectives, z_1 and z_2

Although the Pareto front is of interest, for the design problems described above, we are particular interested in at least one of the end-points of this front. End-points correspond to the solution of individual single criterion problems. We are interested in these end-points because they correspond either to the original design criterion or to one or other design constraints that we have relaxed to gain an understanding of the impact of these criteria. Therefore, although the Pareto front as a whole is of interest, the end-points are particularly relevant. We therefore wish to define a fitness function that emphasises the end-points and hopefully drives the evolutionary algorithm to improve these as much as possible but without sacrificing the Pareto front completely. If the Pareto front were not of interest, we could simply solve a set of single criterion problems.

To achieve the desired fitness values that emphasise not only the Pareto front but especially the end-points of that front, we have defined a rank based fitness function which combines the ranks assigned to each point with respect to each criterion individually:

$$f = 1 - \frac{I_1 \odot I_2 \odot \ldots \odot I_{n_c}}{n_p^{n_c}}$$

where $\odot$ represents the element-wise or Hadamard product of two vectors. The vectors I_j, $j = 1, \ldots, n_c$, each of length n_p, are the indices for each point of their position when the points are sorted with respect to criterion j. If two or more points have the same objective function value, they are implicitly coalesced prior to the assignment of rank for that objective function and hence given the same ranking. The product of the individual rankings denotes the fitness $f_i \in [0,1)$ for $i = 1,\ldots,n_p$ where larger values indicate better fitness. n_p is the number of points and n_c the number of criteria.

Table 1.1 Points and their fitness for the illustration example in Fig. 1.1 sorted in decreasing order according to the rank based fitness

z_1	z_2	I_1	I_2	$I_1 \odot I_2$	f
5	1	6	1	6	0.93
1	8	1	7	7	0.91
3	3	3	3	9	0.89
7	2	8	2	16	0.80
2	9	2	8	16	0.80
4	6	4	5	20	0.75
6	4	7	4	28	0.65
5	7	5	6	30	0.63
9	9	9	9	81	0.00

Table 1.1 illustrates the values of the various vectors for the points shown in Fig. 1.1 sorted according to the fitness value assigned to the points. The best fitness values are for the two extreme points, (5,1) and (1,8) with the next best point being the remaining non-dominated point, (3,3).

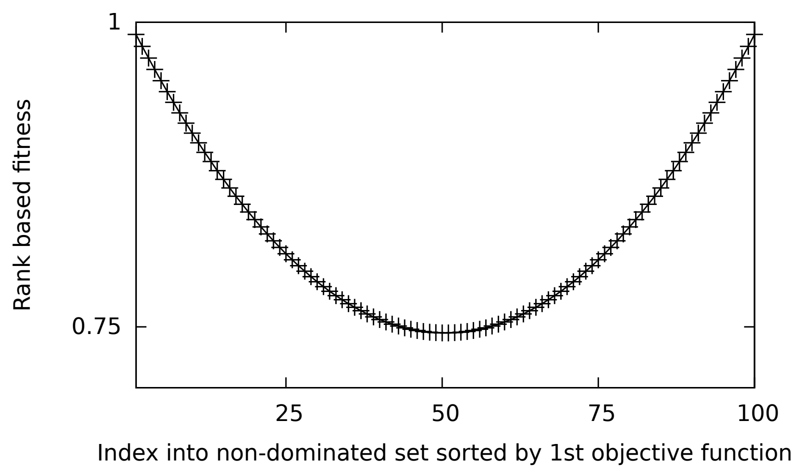

Fig. 1.2 The fitness of solutions along the Pareto front if all solutions are non-dominated. In the plot, the solutions have been sorted according to the first objective function

In general, the largest fitness value achievable is

$$1 - \frac{1}{n_p^{n_c}}$$

which approaches 1 asymptotically as $n_p \to \infty$. This value can only be achieved if one point dominates all the rest. The lowest fitness value is 0, illustrated in the table by the last row, a point that is worst in both criteria. For a population that consists solely of non-dominated points, the fitness along the Pareto front would start at a maximum value and decrease until the middle point of the set is reached and would then start increasing again. The maximum fitness achievable, in this case, will be

$$f\text{max} = 1 - \frac{1}{n_p^{n_c-1}}$$

and the minimum fitness

$$f_{\text{min}} = 1 - \frac{1}{4n_p^{n_c-2}} \quad .$$

With $n_c = 2$ and $n_p = 100$, $f\text{max} = 0.99$ and $f_{\text{min}} = 0.75$. This is illustrated in Fig. 1.2.

1.3 A Multi-objective Plant Propagation Algorithm

The fitness function defined in the previous section could be used with most population based evolutionary algorithms, such as a genetic algorithm. However, we have had good experience with the *Strawberry* algorithm (Salhi and Fraga, 2011), an implementation of a *plant propagation* nature inspired evolutionary method. Further evidence of the power of this approach has been provided by Merrikh-Bayat (2015).

The Strawberry algorithm was originally implemented for single criterion optimisation so it has been extended here for multi-objective problems. This extended algorithm is shown in Algorithm 1. The basic premise is that plants that are in a good position (fertile soil, plenty of water) will reproduce with greater probability but will tend to do so in the vicinity of where they are. Less often, plants which are not well situated will reproduce through longer distance methods. In the Strawberry algorithm, both single and multiple objective versions, each member of the population can generate a number of runners, proportional to that member's fitness, to define new points a distance away proportional to 1 minus the fitness, with all values randomly chosen.

1.4 Case Study: Off-grid Energy Systems Design with Renewable Energy

Mining operations often are located in geographically remote regions of the planet. These operations are seldom connected to grid supplies of energy, either electrical or fuel. As a result, the operations typically require transport of fuel, e.g. diesel,

Algorithm 1 The Strawberry plant propagation algorithm (Salhi and Fraga, 2011) extended for multi-objective optimisation

Given: $f(x)$, a vector function; n_g, number of generations to perform, n_p, the propagation size; n_r, maximum number of runners to propagate.
Output: z, vector approximation to Pareto front.

```
p ← initial random population of size n_p
for n_g generations do
    prune population p, removing similar solutions
    N ← fitness(p)                                  ▷ Use rank based fitness
    p̃ ← φ                                           ▷ Empty set
    for i ← 0...n_p do
        x ← select(p,N)                             ▷ Tournament fitness based selection
        for each runner to generate do              ▷ Number proportional to fitness rounded up
            x̃ ← new solution(x, 1 − N)              ▷ Distance inversely proportional to fitness
            p̃ ← x̃ ∪ p̃                               ▷ Add to new population
        end for
        p ← p\x                                     ▷ Remove from old population
    end for
    p ← p̃ ∪ Nondominated(p)                         ▷ New population with elitism
end for
z ← Nondominated(p)
```

over large distances using trucks or equivalent. There is a desire to reduce the need for the transport of fuel and one possibility is the use of local energy generation. This local generation can be one of solar photovoltaic (PV), solar thermal or wind turbines or a combination of these. Local generation may or may not have economic benefits but such generation will usually have a positive environmental impact when compared with transporting fuel.

Beyond the economics of the choice of generation technology, whether to generate locally at all is the key issue of continuous operation. Many mining sites operate 24 hours a day. A distinguishing factor of the renewable energy generation technologies mentioned above is that they have variable output. Solar based technologies obviously do not generate energy when the sun is not in the sky. Wind turbines can generate energy through day and night but the amount will vary from one moment to the next.

The variability of generation requires changing the design of the mining operation to incorporate energy storage. A number of storage options can be considered for large scale operations: molten salts, pumped hydraulic and compressed air. The design problem then requires identifying the appropriate combination of both generating and storage technologies to minimise the cost of the mining operation. We have previously addressed the minimisation of capital cost (Amusat et al, 2015b).

The optimisation problem for the design problem included a constraint that specified that the power and heat demands of the mining operation had to be satisfied fully from local generated energy from renewable sources. However, in practice, due to the variability of the energy supplies, even with storage, designing for complete reliance on local generation will lead to over-design. Instead, it is more appropriate to design for almost complete reliance on local generation but allowing for the use of fuel brought in from off-site. In other words, from an optimisation point of view, it may be useful to relax the constraint and gain insight into how much impact allowing the use of off-site energy sources, hopefully infrequently, may have.

An analysis of the impact of the constraint was undertaken using a scenario based approach (Amusat et al, 2015a). In this approach, a set of scenarios was generated, with each having a different solar profile over the period of time considered. Each profile was generated randomly. The single objective optimisation problem was solved for each scenario and the resulting design analysed in terms of how likely it was to not meet the demand under different solar profiles. Although this approach was useful, it did not necessarily represent the best solutions possible in a probabilistic sense. A more rigorous approach would be to consider relaxing the demand satisfaction constraint and treating the design problem as a bi-criteria optimisation problem, as discussed above. The demand constraint is a desired attribute of a design but not a hard constraint.

Using the models developed by Amusat et al (2015a) and adding the probability of not meeting demand as a second objective, the problem is now

$$\min_d z = \begin{cases} c(d) \\ p(d) \end{cases} \tag{1.1}$$

where d are the design variables, $c(\cdot)$ is the capital cost of the energy generation and storage systems and $p(\cdot)$ is the probability of not meeting the demand fully with local generation sources. The probability is 1 minus the reliability. Reliability is a measure of the ability of an energy system to deliver power to all points of consumption with the frequency, the duration and the extent required by the operation (Osborn and Kawann, 2001).

The evaluation of the objective functions proceeds as follows:

1. For a given d, the generation and storage technologies are defined, resulting in an energy system for the mining operation.
2. The resulting design is then evaluated over a number of randomly generated scenarios based on a probability distribution function describing the variability of solar irradiance for each hour of each day in the period of operation.
3. The probability of not meeting demand is simply the ratio of the number of scenarios where the demand was not met for design d, for at least one time period, and the number of total scenarios. A value of 0 means that the design is able to meet the demand under all likely solar conditions; a value of 1 means that the design never fully meets the demand, always requiring the import of off-site fuel for at least one time period over the full duration for each scenario.

As a starting point, we have solved this bi-criteria problem of dimension 8 using NSGA-II (Deb, 2000) with population size 100, 150 generations, crossover rate of 0.25, mutation rate of 0.25, with binary tournament selection, intermediate crossover and Gaussian mutation. The non-dominated objective function values for 5 attempts are presented in Fig. 1.3. At the scale used, NSGA-II appears to identify the set of non-dominated designs well. From an engineering point of view, we do see that the designs are not that sensitive to the variability in solar irradiance. This is not entirely surprising as the case study consists of the Collahuasi mine located in the Atacama region of Chile where most days have completely clear skies and so the irradiance is relatively constant and predictable.

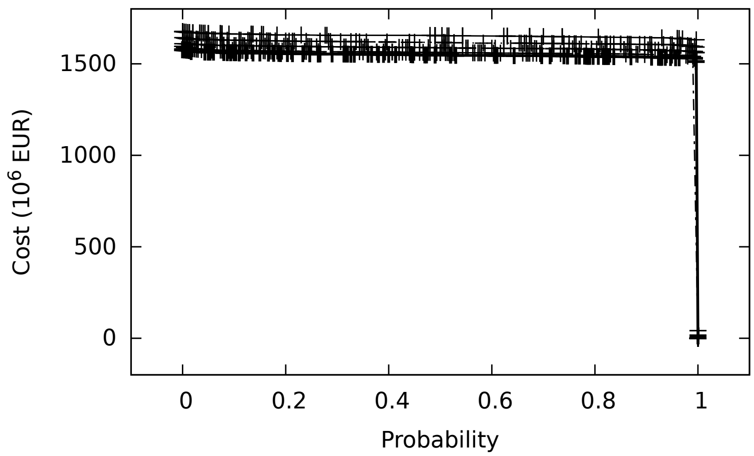

Fig. 1.3 Final sets of non-dominated points resulting from 5 attempts at the off-grid design problem using NSGA-II (Deb, 2000) where the probability is the likelihood of not satisfying the energy demands of the mining operation with only local energy generation

Figure 1.4 shows the different sets of non-dominated points for the left side of the plot shown in Fig. 1.3. It is this part of the plot we are most interested in as the designs here are those that will not very often require off-site fuel. In zooming in, we see that the solutions obtained are similar in objective function value from one attempt with NSGA-II to another.

For comparison, we have also applied the Strawberry method using the rank based fitness function described above with population size 100, 150 generations, and $n_r = 5$. The aim of this fitness function is to emphasise the end-points of the Pareto front and, due to the asymmetry in evaluation, particularly the left end-point. Figure 1.5 shows the resulting solutions. Of note,

1. The Strawberry algorithm is less consistent over the full range of probability values.
2. The number of points in the set of non-dominated points is small compared with the sets generated by NSGA-II.
3. The distribution of points is less even than it is for the NSGA-II case with points concentrated more towards the extremes of the Pareto front than towards the centre. Note that the use of diversity control reduces the number of points at the right extreme.
4. The cost objective function values obtained are lower than those obtained using NSGA-II.

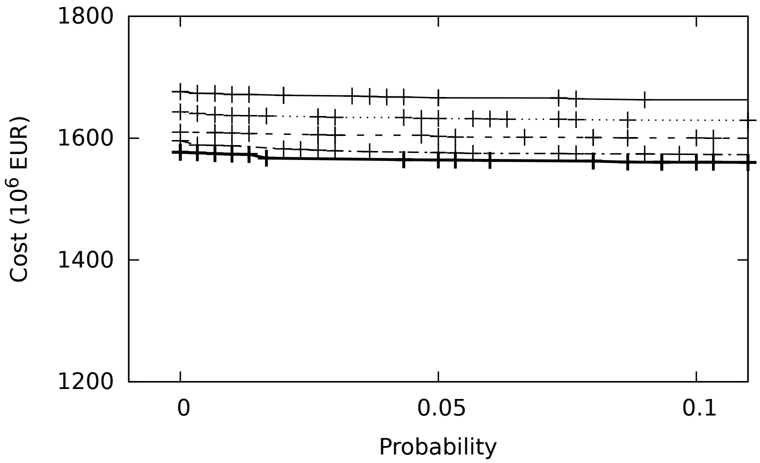

Fig. 1.4 A zoomed in view of Fig. 1.3

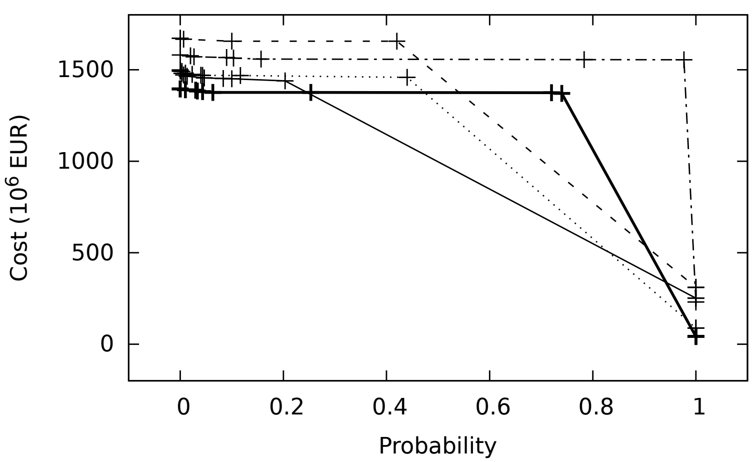

Fig. 1.5 Final sets of non-dominated points resulting from 5 attempts at the off-grid design problem using the new multi-objective Strawberry algorithm where the probability is the likelihood of not satisfying the energy demands of the mining operation with only local energy generation

If we zoom in as we did for the NSGA-II results, we see (Fig. 1.6) that the cost objective function is often better than what is obtained with NSGA-II.

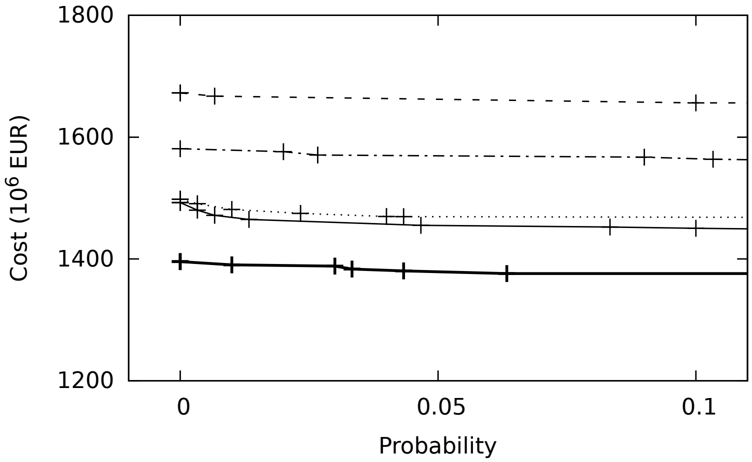

Fig. 1.6 A zoomed in view of Fig. 1.5

1.4.1 Analysis of the designs

The aim of using a multi-objective approach to solve the constrained single objective design problem is to gain insight into the impact of the particular constraint on the designs obtained. In the off-grid design problem, the question the design engineer would like to answer is: Does the constraint of not using any off-site fuel lead to a significant over-design?

Table 1.2 Design points with low probability of not meeting the demand for the mining operation. The first row is a design which should **always** meet the demand and the second row is one in which the demand will be met all but 1% of the possible solar profiles that could be encountered. The *G* columns are generation (power tower, PT, and photovoltaic, PV). The *S* columns refer to storage: MS for molten salts, PH for pumped hydro and CA for compressed air. Finally, *E* columns indicate the peak electricity release rate from storage.

	G_{PT} (MW)	S_{MS} (MWh)	E_{MS} (MW)	G_{PV} (MW)	S_{PS} (MWh)	E_{PS} (MW)	S_{CA} (MWh)	E_{CA} (MW)	p	c 10^6 €
	1257	6022	180	0.00	2746	89.2	0.00	60.26	0.000	1396
	1246	6036	177	0.00	2740	94.8	0.00	64.80	0.010	1390
	1238	6000	179	0.00	2749	96.4	25.98	64.73	0.030	1388
	1235	6021	179	0.87	2839	93.5	1.44	58.34	0.033	1383
	1232	6029	177	1.54	2795	93.6	0.00	59.79	0.043	1381
	1228	6021	179	0.00	2736	89.8	0.00	65.10	0.063	1376
Δ (%)	2	1	2	100	4	7	100	10	100	1

Table 1.2 shows the values of the first 6 design points, counting from the left, from the bottom graph in Fig. 1.6. The cost does not include the actual cost of the off-site fuel and the cost of transporting that fuel so this table (and the results

discussed earlier) only allows us to analyse the impact on the physical structure of the mining operation's energy systems.

The final row shows the variation as a percentage of the maximum value for each design variable and objective function. From the first design through to the last one, there is a change of 1% in cost. Some variables are stable whereas the use of photovoltaic generation and compressed air storage have large changes: the feasible design does not include either of these technologies. Introducing them allows the cost to be reduced but at the expense of not meeting demands in all scenarios.

For the design engineer, the main conclusion is that the feasible design is not over-specified. In fact, some of the slightly less expensive designs introduce more complexity by the incorporation of further alternative technologies. Generally, the simpler the design, the more attractive it is so the feasible design is favoured even more.

The ability to perform this analysis enables the engineer to have the confidence necessary to move to the next stage of design: the detailed specification of the individual technologies and further modelling, simulation and optimisation.

1.5 Conclusions

The use of multi-objective optimisation can provide useful insight into the impact of constraints on designs. By converting a constraint to an extra objective, the approximation of the Pareto front for the design problem will help determine, for instance, whether the "feasible" design is over-constrained or not. To ensure that the impact of the relaxation of a constraint is understood, it is necessary to have a good approximation to the Pareto front at the extremes. This motivates the definition of a fitness function to provide the appropriate selection pressure for evolutionary methods.

A rank based fitness function has been presented. An implementation has been incorporated in a new multi-objective plant propagation algorithm based on the existing single objective *Strawberry* algorithm (Salhi and Fraga, 2011). The procedure has been applied to a problem in off-grid operation of large scale mining where there is a desire to reduce the cost and environmental impact of using fuel brought in from a long distance away. The results demonstrate the effectiveness of both the fitness function and the multi-objective Strawberry method.

1.5.1 Acknowledgements

The authors would like to acknowledge the funding provided by the Nigerian government through the Presidential Scholarship for Innovation and Development (PRESSID) scheme.

References

Amusat O, Shearing P, Fraga ES (2015a) Optimal integrated energy systems design incorporating variable renewable energy sources. In: Roskilly T, Rajendran K, Ling-Chin J (eds) Proceedings of Sustainable Thermal Energy Management (SusTEM), pp 245–253

Amusat O, Shearing P, Fraga ES (2015b) System design of renewable energy generation and storage alternatives for large scale continuous processes. In: Gernaey KV, Huusom JK, Gani R (eds) 12th International Symposium on Process Systems Engineering and 25th European Symposium on Computer Aided Process Engineering, Elsevier B. V., pp 2279–2284

Coello Coello CA, Landa Becerra R (2009) Evolutionary multiobjective optimization in materials science and engineering. Materials & Manufacturing Processes 24(2):119–129

Deb K (2000) An efficient constraint handling method for genetic algorithms. Comput Methods Appl Mech Engng 186:311–338

Fiandaca G, Fraga ES, Brandani S (2009) A multi-objective genetic algorithm for the design of pressure swing adsorption. Engineering Optimization 41(9):833–854

Merrikh-Bayat F (2015) The runner-root algorithm: A metaheuristic for solving unimodal and multimodal optimization problems inspired by runners and roots of plants in nature. Applied Soft Computing 33:292–303

Osborn J, Kawann C (2001) Reliability of the us electricity system: Recent trends and current issues. Tech. rep., Energy Analysis Department, Ernest Orlando Lawrence Berkeley National Laboratory, LBNL-47043, Berkeley, CA

Pardalos PM, Žilinskas A, Žilinskas J (2016) Non-convex Multi-objective Optimization. Springer

Salhi A, Fraga ES (2011) Nature-inspired optimisation approaches and the new plant propagation algorithm. In: Proceedings of ICeMATH 2011, The International Conference on Numerical Analysis and Optimization, Yogyakarta, pp K2:1–8

Törn A, Žilinskas A (1989) Global Optimization. Springer-Verlag, Berlin