Proceedings of the 2015 Winter Simulation Conference
L. Yilmaz, W. K V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, eds.

A NEW MYOPIC SEQUENTIAL SAMPLING ALGORITHM FOR MULTI-OBJECTIVE PROBLEMS

Juergen Branke

Warwick Business School
The University of Warwick
Coventry, CV4 7AL, UK

Wen Zhang

Warwick Business School
The University of Warwick
Coventry, CV4 7AL, UK

ABSTRACT

In this paper, we consider the problem of efficiently identifying the Pareto optimal designs out of a given set of alternatives, for the case where alternatives are evaluated on multiple stochastic criteria, and the performance of an alternative can only be estimated via sampling. We propose a simple myopic budget allocation algorithm based on the idea of small-sample procedures. Initial empirical tests show encouraging results.

1 INTRODUCTION

A typical use of stochastic simulation models is to identify the best out of a given set of alternatives, where best is defined in terms of each alternative's unknown expected value, and the expected values have to be inferred through statistical sampling (where each simulation run constitutes a sample). How to do this efficiently is the question addressed in the area of Ranking and Selection (R&S), an area that has seen substantial interest in recent years. The state-of-the-art procedures allocate the sampling budget sequentially onto the different alternatives, based on the observations made so far. A comparison of different sequential R&S techniques can be found in (Branke, Chick, and Schmidt 2007).

This paper considers the case where alternatives are evaluated along more than one stochastic criteria. In this case, what constitutes best is no longer uniquely defined, but there are usually several equally good solutions with different trade-offs between the objectives. The goal then becomes to identify all Pareto optimal solutions (the so called Pareto set), i.e., all solutions that are not dominated by any other alternative. An alternative is said to dominate another solution if it is at least equally good in each criterion, and strictly better in at least one criterion.

In this paper, we propose an extension of the small sample procedures proposed in Chick, Branke, and Schmidt (2010) to the multi-objective case. By being myopic and only estimating the benefit of allocating a few additional samples to one alternative, small sample procedures can avoid various asymptotic approximations. More specifically, in each iteration of sample allocation, we will only allocate samples to one alternative, which is the alternative that has the highest probability of changing the observed Pareto set.

So far, there are only few papers on multi-objective ranking & selection (MORS) problems. The most prominent method is probably the multi-objective optimal computing budget allocation (MOBCA) method proposed by Lee, Chew, Teng, and Goldsman (2010). MOCBA is the multi-objective version of the Optimal Computing Budget Allocation (OCBA) algorithm (Chen, Lin, Yücesan, and Chick 2000) and has been derived by minimizing the upper bounds for Type I and Type II errors asymptotically (for infinite sampling budgets). At each iteration, the probability of Type I and Type II error are estimated, and an allocation rule is applied that aims at minimizing the higher of the two. Calculating the allocation of samples at each stage is an iterative process. For ease of computation, Chen and Lee (2010) also proposed a simplified version of MOCBA by ignoring the impact of the allocation quantity on the role of domination, and using the variance associated with each design directly in computing the allocation quantity. MOCBA has been extended by Teng, Lee, and Chew (2010) to be able to take into account an indifference zone, and adapted by Lee, Chew, and Teng (2010) for different measures of selection quality such as expected opportunity cost.

Yan, Zhou, and Chen (2010) consider the problem of identifying the subset of m simplest (lowest complexity) and good enough designs among a larger set of designs. Since complexity here is a deterministic criterion, this constitutes a 2-objective problem with one stochastic (quality) and one deterministic (complexity) objective. They also propose an adaptation of OCBA.

Besides OCBA, researchers also adapted other R&S methods from the single-objective to the multi-objective case. For instance, Zhang, Georgiopoulos, and Anagnostopoulos (2013) presented an S-Race algorithm which attempts to eliminate alternatives as soon as there is sufficient statistical evidence of their inferiority relative to other alternatives with respect to all objectives. Experimental results implied that S-Race typically offers cogent computational advantages in exchange for relatively small discrepancies in the set of Pareto-optimal models (Zhang, Georgiopoulos, and Anagnostopoulos 2013). Another multi-objective racing algorithm to be used to select individuals within evolutionary algorithms has been proposed by Marceau-Caron and Schoenauer (2014). They proposed to directly estimate the probability of each individual to survive to the next generation, and use Hoeffding races to drop alternatives early. Yahyaa, Drugan, and Manderick (2014) extended the knowledge gradient (KG) polices to the multi-objective, multi-armed bandits problem to efficiently explore the Pareto optimal arms and found that KG outperformed an upper confidence bound policy. Kabirian and Olafsson (2009) proposed a new heuristic iterative algorithm for finding the best design when there are some stochastic constraints and showed the effectiveness of the algorithm in dealing with different test problems in numerical results.

The paper is organized as follows. Section 2 formalizes the problem and describes the assumptions. Section 3 describes the proposed small-sample multi-objective optimization procedure. Section 4 presents empirical simulation results, and the paper concludes in Section 5 with a summary and some suggestions for future work.

2 PROBLEM FORMULATION

In this section, the problem of multi-objective R&S is introduced and the notation is specified. Given H objectives and a set of m designs with the true unknown performance of each design i in objective h veing denoted by w_{ih}. Assuming minimization throughout this paper, a design i is said to dominate design j $(i \succ j)$ if and only if $w_{ih} \leq w_{jh}$ for all objectives and $w_{ih} < w_{jh}$ for at least one objective. The goal is to identify all Pareto-optimal designs (the Pareto set), i.e., all designs that are not dominated by any other design.

The performances of each design in each objective need to be estimated via sampling (e.g., running a stochastic simulation model). Let $\mathbf{X}_i$ be a matrix that contains the simulation output for design i. Then $\mathbf{X}_i$=(X_{ihn}), where X_{ihn} is the h-th objective of design i for simulation replication n. Let furthermore w_{ih} and σ_{ih}^2 be the unknown mean and variance of alternative i, which can only be estimated using the simulation outputs X_{ihn}. We assume that

$$\{X_{ihn} : n = 1,2,...\} \overset{iid}{\sim} \mathcal{N}(w_{ih}, \sigma_{ih}^2), \text{for } i = 1,2,...,m \text{ and } h = 1,2,...H.$$

Let n_i be the number of samples taken for alternative i so far, $\bar{x}_{ih}$ be the sample mean and $\hat{\sigma}_{ih}^2$ be the sample variance. Then, we will get an observed Pareto set based on the $n = \sum_i n_i$ simulations so far. As n_i increases, $\bar{x}_{ih}$ and $\hat{\sigma}_i^2$ will be updated and the observed Pareto front may change accordingly.

If alternative i is to receive another τ_i samples, and $\bar{y}_{ih}$ is the average of the new samples in objective h, then the new overall sample mean in each objective can be calculated as

$$z_{ih} = \frac{n_i \bar{x}_{ih} + \tau_i \bar{y}_{ih}}{n_i + \tau_i}. \tag{1}$$

Before the new samples are observed, the sample average that will arise after sampling, denoted as Z_{ik}, is a random variable, and we can use the predictive distribution for the new samples (DeGroot 2005) and get

$$Z_{ih} \sim St(\bar{x}_{ih}, n_i * (n_i + \tau_i)/(\tau_i * \hat{\sigma}_{ih}^2), n_i - 1)$$

where $St(\mu, \kappa, \nu)$ denotes the student distribution with mean μ, precision κ and ν degrees of freedom.

Chick, Branke, and Schmidt (2010) use different methods to define the objective of a selection procedure. In this paper, the simplest method of Probability of Correct Selection (P{CS}) is adopted, but it should be straightforward to extend the basic ideas also to other objectives such as expected opportunity cost (linear loss). Specifically, we will say a selection is correct when the selected Pareto set is the true Pareto set, and wrong otherwise. For N replications of an experiment, the observed $P\{CS\}$ can be calculated as

$$P\{CS\} = \left(\frac{N_c}{N}\right)$$

where N_c is the number of replications for which the method returned the correct Pareto set.

3 MYOPIC MULTI-OBJECTIVE BUDGET ALLOCATION PROCEDURE

Inspired by the idea of the small-sample EVI procedure proposed by Chick, Branke, and Schmidt (2010), we propose a new myopic budget allocation rule for multi-objective scenarios called Myopic Multi-Objective Budget Allocation (M-MOBA). The small sample EVI procedure means that in any given stage, instead of allocating samples to each alternative, we will only allocate $\tau_j > 0$ samples to one alternative j and τ_i=0 for all other alternatives. If the additional samples taken do not change our decision as to which alternatives belong to the Pareto set, then, in the myopic sense of only looking one step ahead, they did not have any value. On the other hand, if the additional samples taken cause us to regard a different set as Pareto set, then this information is assumed to be helpful.

Following this logic, we only allocate simulation budget to the alternative that has the highest probability to change our perception as to which solutions belong to the Pareto set at the next simulation stage. For the sake of convenience, the allocation rule is explained by assuming that there are two objectives for each alternative so that the Pareto set and the dominance relationship can be visualized in a two-dimensional coordinate system. However, extension to more than two objectives should be straightforward. In the following subsections, we will explain how to calculate the probability of changing the Pareto set by sampling a particular alternative M, in different situations. In all cases, we assume we already allocated $n = \sum_i n_i$ samples.

3.1 Alternative M is in the Pareto Set

For the point on the observed Pareto front, there are two situations.

3.1.1 M solely dominates another solution

Consider the situation depicted in Figure 1, with the observed Pareto front as highlighted by the dashed line composed of points A, M, C, D. Note that alternative M is the sole solution that dominates alternatives B, E, F, G, and alternatives B, E, F would become non-dominated if M was removed. Meanwhile, alternative B is the 'nearest' neighbor of M in the x axis direction and alternative F is the 'nearest' neighbor of M in the y axis direction.

We want to calculate the probability that the current Pareto front will change if we allocate τ additional simulation samples to M. Obviously, if we only allocate samples to M, all other solutions can be considered deterministic. Then, the Pareto set changes if and only if the new mean estimate for solution M after sampling falls outside the shaded area in Figure 2 and either dominates one of A, C, D or becomes dominated itself. In the rest of the paper, for the sake of convenience, we assume alternative B is the 'nearest' neighbor of M in both dimension and only B is included in the pictures to explain the method.

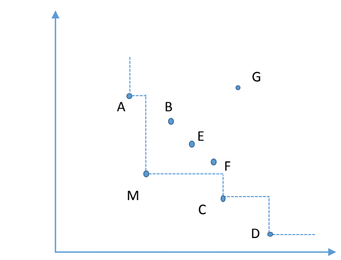

Figure 1: An observed Pareto front

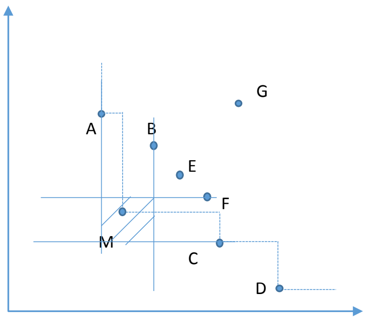

Figure 2: The Pareto set will change if and only if the estimated mean of alternative M will fall outside the shaded area

Since we assume that the samples in the two objectives are independent, we can calculate the probability for M to remain in the shaded area separately for each objective, and multiply them to get the probability that the 'new' M falls into the shaded area.

Specifically, we assume the four vertices of the shaded area are a, b, c, d and the coordinate of each vertex is $a(x_1,y_2)$, $b(x_1,y_1)$, $c(x_2,y_2)$ and $d(x_2,y_1)$ accordingly as shown in Figure 3. Hence, the probability P that X falls into the shaded area, can be calculated as

$$P = P_x * P_y$$

where $P_x = P\{x_1 < X_{x,n+\tau} < x_2\}$ and $P_y = P\{y_1 < X_{y,n+\tau} < y_2\}$.

The values of x_1, x_2, y_1 and y_2 are easily identified. Since we want to analyze the effect caused by the movement of point M, we first remove M and get the Pareto front composed of the remaining points. In the example, the new Pareto front is composed of points A, B, C and D, see Figure 4. This Pareto front is denoted as F_2. Then, x_1 and x_2 are the x-coordinate values of the 'nearest neighbor points' around point X on the x-axis which belong to F_2, and y_1 and y_2 are the y-coordinate value of the 'nearest neighbor points' around point X on the y-axis that belong to F_2. In this case, x_1 is the x-coordinate value of point A, x_2 is the x-coordinate value of point B, y_1 is the y-coordinate value of point C and y_2 is the y-coordinate value of point B.

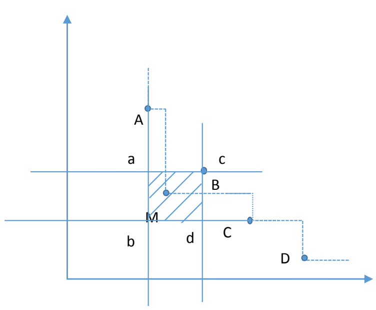

Figure 3: Vertex of shaded area

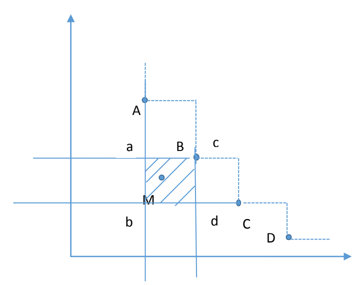

Figure 4: The Pareto set will change if and only if the estimated mean of alternative M will fall outside the shaded area

3.1.2 M Does Not Solely Dominate Another Alternative

In the situation depicted in Figure 5, alternative M does not solely dominate another alternative that would become non-dominated if M was removed. Thus, the Pareto set changes if and only if the new samples will lead to M dominating a previously non-dominated solution, or M becoming dominated itself. The shaded area in Figure 6 shows the locations for which there is no change in the Pareto set if M falls into this area after the additional samples.

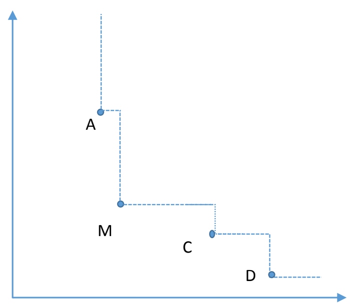

Figure 5: M does not solely dominate another alternative

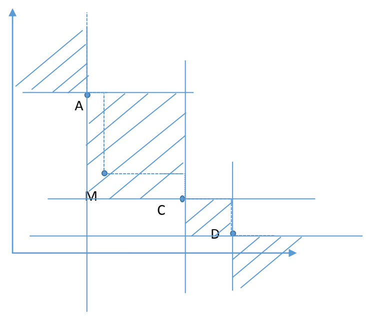

Figure 6: The Pareto set will change if and only if the estimated mean of alternative M will fall outside the shaded area

This area is less tricky to calculate compared to the previous one since the vertex of each sub area is only decided by the points on F_2 and the process of finding 'nearest neighbor points' is omitted, but several probabilities (rectangles) need to be added up.

3.2 M is Not in the Pareto Set

In Figure 7, solution M is not in the Pareto front (dashed line composed of points A, B, C, D). In this case, the Pareto set changes if and only if solution M becomes non-dominated, i.e., falls outside the shaded region in Figure 8, which can be calculated in a straightforward way.

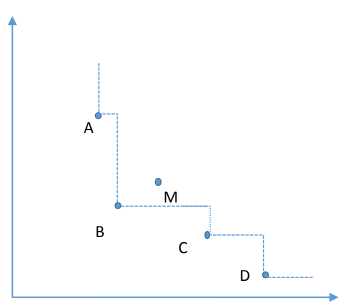

Figure 7: Observed Pareto front with M dominated

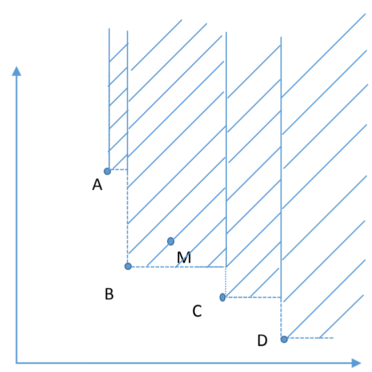

Figure 8: The Pareto set will change if and only if the estimated mean of alternative M will fall outside the shaded area

3.3 M-MOBA Algorithm

Based on the above analysis, we can formulate the small-sample multi-objective budget allocation procedure as follows.

Procedure M-MOBA

1. Specify a first-stage sample size $n_0 \geq 5$, and a number of samples $\tau > 0$ to allocate per subsequent stage. Specify stopping rule parameters.
2. Sample X_{ihn}, $i = 1, \ldots, m; h = 1, \ldots, H; n = 1, \ldots, n_0$ independently, and initialize the number of samples $n_i \leftarrow n_0$
3. Determine the sample statistics $\bar{x}_{ih}$ and $\hat{\sigma}^2_{ih}$, and the observed Pareto front.
4. WHILE stopping rule not satisfied DO another stage:
   (a) For each alternative i, calculate the probability P_i that the new samples will lead to a change in the Pareto set.
   (b) Allocate τ samples to the alternative that has the largest P_i.
   (c) Update sample statistics n_i, $\bar{x}_{ik}$ and $\hat{\sigma}^2_{ik}$ and observe a new Pareto set.
5. Select alternatives on the observed Pareto set.

4 EMPIRICAL SIMULATION RESULTS

In this section, we present empirical experiments using the proposed M-MOBA method and compare its performance with that of equal allocation which simply allocates an equal number of samples to each alternative, and the MOCBA method used by Chen and Lee (2010). Please note that for MOCBA, we did not have accurate performance data, but instead have taken its results from a figure provided in Chen and Lee (2010), so small inaccuracies were unavoidable.

We test two problems taken from (Chen and Lee 2010). We replicate each run using Matlab 1000 times and calculate $P\{CS\}$ based on the results. Stopping rule used in the experiments is the sampling budget rule, i.e., continuing sampling up to a specified available total budget.

4.1 Three Designs Problem

In this benchmark problem, there are 3 designs and each of them is evaluated according to 2 objectives. Parameters of the designs are shown in Table1.

Table 1: 3 designs.

Index	Obj. 1	Std. dev. 1	Obj. 2	Std. dev. 2
0	1	5	2	5
1	3	5	1	5
2	5	5	5	5

In the experiment, the initial number of samples for each design is n_0=5, the total computing budget is 500 and the number of samples allocated in each iteration is $\tau = 1$.

The resulting $P\{CS\}$ over the budget allocated is shown in Figure 9. The results show that the probability of correct selection using our algorithm is generally higher than equal allocation with the same simulation budget. Also, $P\{CS\}$ using the M-MOBA converges faster to 1 than equal allocation. M-MOBA performs very similar to MOCBA on this problem.

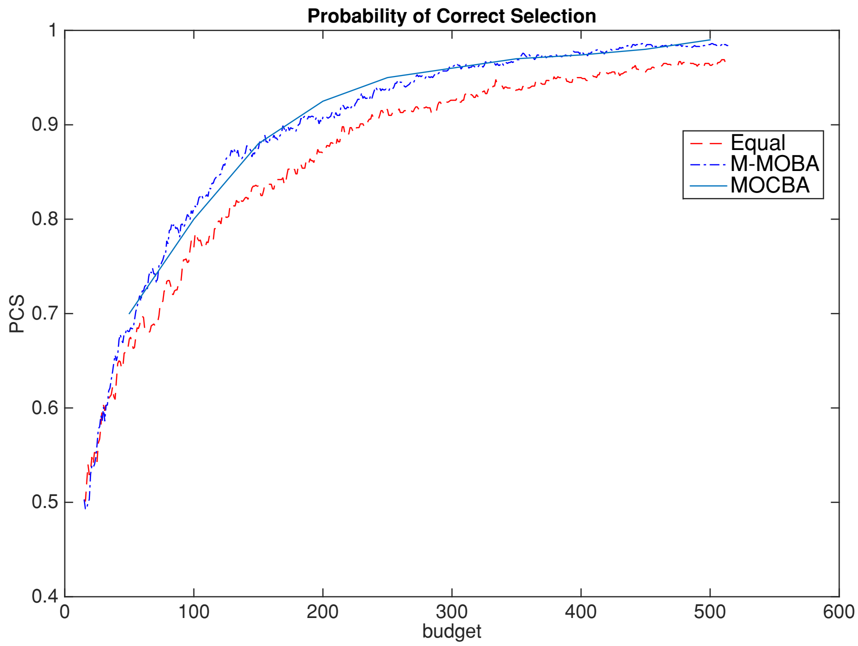

Figure 9: Comparison of $P\{CS\}$ for different algorithms in the 3-design case.

4.2 Sixteen Designs Problem

In this benchmark problem, there are 16 designs, each of them is evaluated according to 2 objectives, and the standard deviation of each alternative in each objective is set to 2. Expected values of each design are shown in Table 2.

Table 2: Benchmark problem with 16 alternatives and two objectives. Standard deviation for all designs is 2 in each objective.

Index	Obj. 1	Obj. 2
0	0.5	5.5
1	1.9	4.2
2	2.8	3.3
3	3	3
4	3.9	2.1
5	4.3	1.8
6	4.6	1.5
7	3.8	6.3
8	4.8	5.5
9	5.2	5
10	5.9	4.1
11	6.3	3.8
12	6.7	7.2
13	7	7
14	7.9	6.1
15	9	9

Again, we conduct 1000 independent replications with $n_0 = 5, \tau = 1$ and a total computing budget of 4000 for M-MOBA and equal allocation, and read the performance values for MOCBA from a figure in the book by Chen and Lee (2010).

Initial results of M-MOBA were good for small budgets, but then the algorithm got stuck and stopped improving after having allocated approximately 1000 samples. We suspected the reason to be due to the limited precision of the computations. And indeed, a closer examination revealed that as the budget becomes larger and larger, the probability that an additional sample would influence the Pareto set becomes very small and Matlab can no longer differentiate between the alternatives and allocates samples only to the first alternative which obviously brings the algorithm to a halt. As a simple fix to this problem, we then tried two slight modifications. First, we assumed we allocate $\tau = 10$ samples instead of only one sample to an alternative when calculating the probability of observing a change in the Pareto set. Obviously, the impact of allocating ten samples is larger than allocating only one sample, and the probabilities are thus larger as well. However, we only use $\tau = 10$ for calculating the probabilities, and still only allocate a single sample $\tau = 1$ to the design for which we predict the highest probability of change. Second, if we still get zero probability of change for all alternatives, then we allocate samples equally among them. Results are summarized in Figure 10, where the modified M-MOBA algorithm is labeled M-MOBA-mod. As can be seen, this simple modification solves the problem, and the PCS of the algorithm keeps improving.

In comparison to MOCBA, it seems that M-MOBA is equivalent for small budgets, but then becomes slightly worse as the budget is increased. This may be expected, as M-MOBA is a myopic procedure designed with small budgets in mind, whereas MOCBA is based on asymptotic considerations. It works better in comparison to the simplified MOCBA version proposed by Chen and Lee (2010) over the whole range of budgets, and thus represents a competitive and easy-to-implement alternative to MOCBA.

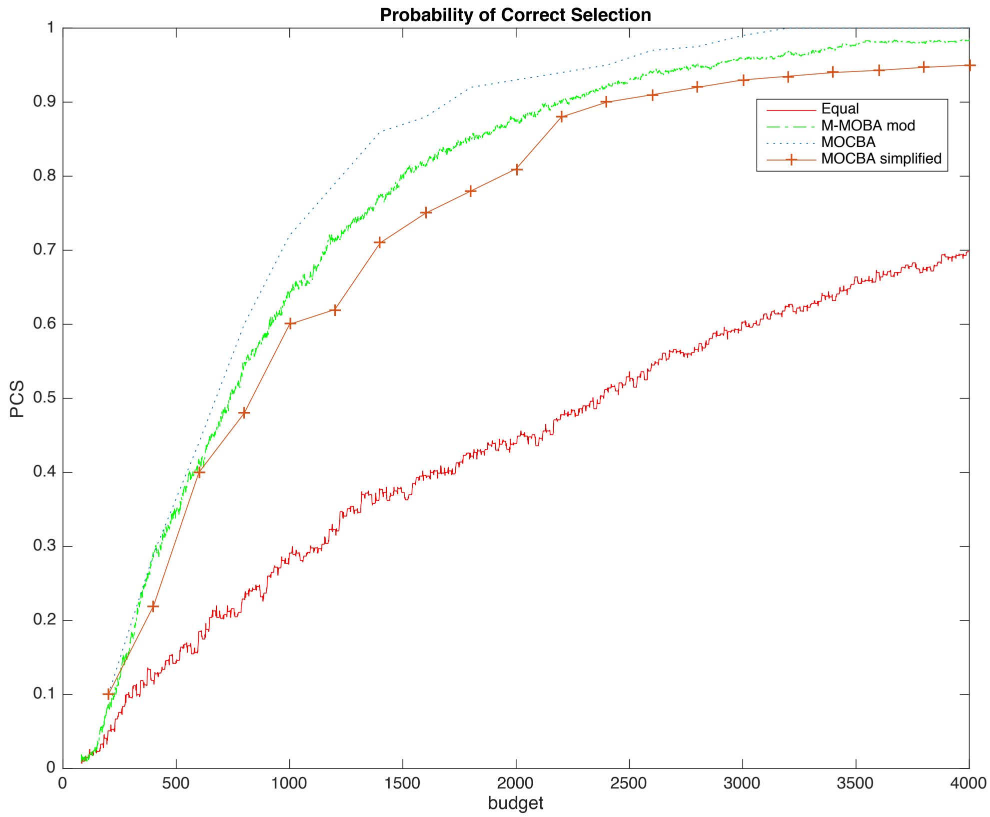

Figure 10: Comparison of $P\{CS\}$ for different algorithms on the 16-design case.

5 CONCLUSION

In this paper, a new simple myopic budget allocation algorithm for solving multi-objective problems has been proposed. This method is an extension of the small-sample EVI procedure proposed by Chick, Branke, and Schmidt (2010) to multi-objective problems. Empirical comparisons to Equal allocation and MOCBA show that the new method works very well, especially in situations where the total budget that can be allocated is small. We observed some problems with numerical precision when allocating many samples, however this could be fixed simply by assuming a larger number of samples τ allocated when computing probabilities and switching to equal allocation in case the probability of change drops to zero for several solutions. More sophisticated ways to deal with the problem of numerical precision will be the next step on our agenda. Also, it should be straightforward to extend the methodology to different stopping rules, different performance measures (e.g., expected opportunity cost instead of probability of correct selection), stochastic constraints, the case when objectives may be evaluated separately or objective values are correlated, as well as situations that consider more than two objectives.

ACKNOWLEDGMENTS

We thank Mr Yang Tao, PhD student of Department of Mathematics, Imperial College London, for his useful suggestions.

REFERENCES

Branke, J., S. E. Chick, and C. Schmidt. 2007. "Selecting a selection procedure". *Management Science* 53 (12): 1916–1932.

Chen, C. H., and L. H. Lee. 2010. *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*. World Scientific.

Chen, C.-H., J. Lin, E. Yücesan, and S. Chick. 2000. "Simulation Budget Allocation for Further Enhancing the Efficiency of Ordinal Optimization". *Discrete Event Dynamic Systems: Theory and Applications* 10:251–270.

Chick, S. E., J. Branke, and C. Schmidt. 2010. "Sequential sampling to myopically maximize the expected value of information". *INFORMS Journal on Computing* 22 (1): 71–80.

DeGroot, M. H. 2005. *Optimal Statistical Decisions*, Volume 82. John Wiley & Sons.

Kabirian, A., and S. Olafsson. 2009. "Selection of the best with stochastic constraints". In *Winter Simulation Conference*, 574–583. Winter Simulation Conference.

Lee, L. H., E. P. Chew, and S. Teng. 2010. "Computing budget allocation rules for multi-objective simulation models based on different measures of selection quality". *Automatica* 46 (12): 1935–1950.

Lee, L. H., E. P. Chew, S. Teng, and D. Goldsman. 2010. "Finding the non-dominated Pareto set for multi-objective simulation models". *IIE Transactions* 42 (9): 656–674.

Marceau-Caron, G., and M. Schoenauer. 2014. "Racing Multi-objective selection probabilities". In *Parallel Problem Solving from Nature (PPSN)*, Number 8672 in LNCS, 631–640: Springer.

Teng, S., L.-H. Lee, and E. P. Chew. 2010. "Integration of indifference-zone with multi-objective OCBA". *European Journal of Operational Research* 203:419–429.

Yahyaa, S. Q., M. M. Drugan, and B. Manderick. 2014. "Knowledge gradient for multi-objective multi-armed bandit algorithms". In *International Conference on Agents and Artificial Intelligence (ICAART)*, 74–83.

Yan, S., E. Zhou, and C.-H. Chen. 2010. "Efficient simulation budget allocation for selecting the best set of simplest good enough designs". In *Proceedings of the Winter Simulation Conference*, 1152–1159. Winter Simulation Conference.

Zhang, T., M. Georgiopoulos, and G. C. Anagnostopoulos. 2013. "S-Race: A multi-objective racing algorithm". In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, 1565–1572. ACM.

AUTHOR BIOGRAPHIES

JUERGEN BRANKE is Professor of Operational Research and Systems of Warwick Business School, University of Warwick, UK. He is Associate Editor for the Evolutionary Computation Journal and Area Editor for the Journal of Heuristics. His research interests include metaheuristics, multiobjective optimization and decision making, optimization in the presence of uncertainty, and simulation-based optimization, and he has published over 150 peer-reviewed paper in international journals and conferences. His e-mail address is Juergen.Branke@wbs.ac.uk.

WEN ZHANG is a PhD student of the Operational Research and Management Science Group at Warwick Business School, University of Warwick. Her research interests mainly lie in multi-objective optimization and decision making. Her email address is phd14wz@mail.wbs.ac.uk.