Proceedings of the 2018 Winter Simulation Conference
M. Rabe, A. A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, eds.

PROVABLY IMPROVING THE OPTIMAL COMPUTING BUDGET ALLOCATION ALGORITHM

Di Wu

School of Industrial & Systems Engineering
Georgia Institute of Technology
755 Ferst Drive NW
Atlanta, GA 30332, USA

ABSTRACT

We boost the performance of the Optimal Computing Budget Allocation (OCBA) algorithm, a widely used and studied algorithm for Ranking and Selection (as known as Best Arm Identification) under a fixed budget. The proposed fully sequential algorithms, OCBA+ and OCBAR, are shown to have better performance both theoretically and numerically. Surprisingly, we reveal that in a two-design setting, a constant initial sample size in a family of OCBA-type algorithms (including the original OCBA) only amounts to a sub-exponential or even polynomial convergence rate of the probability of false selection (PFS). In contrast, our algorithms are guaranteed to converge exponentially fast, as is shown by a finite-sample bound on the PFS.

1 INTRODUCTION

In simulation optimization, the study of Ranking and Selection (R&S) mainly focuses on how to efficiently run simulations to identify the best design among a finite number of candidates. R&S has two major formulations. The fixed confidence setting challenges us to achieve certain confidence level using the least possible simulation effort. The fixed budget setting, on the other hand, requires maximizing the probability of selecting the best design using a fixed budget of simulation runs. The focus of this work is on the fixed budget formulation.

In fixed budget R&S, the Optimal Computing Budget Allocation (OCBA) algorithm in Chen et al. (2000) is considered as the one of the most widely used algorithms. The framework of OCBA has been extended to handle many applications, and similar sequential allocation style has been explored in numerous subsequent works. However, to the best of our knowledge, although the asymptotic properties of simple R&S procedures have been investigated, the formal analysis and characterization of sequential OCBA's performance remains an open research problem.

In view of the current void in OCBA's theoretical performance guarantees, we are motivated to push the boundary by performing in-depth analysis on its behavior, and developing insights for improving its performance. Our contributions are outlined as follows.

1. We propose two fully sequential algorithms, OCBA+ and OCBAR, and provide a finite-sample bound on its probability of correct selection (PCS), guaranteeing an exponential convergence rate.
2. It is revealed that for a two-design case, if OCBA, OCBA+ and OCBAR choose a constant initial sample size, then the PCS converges only at a *sub-exponential* (or even *polynomial*) rate.
3. Numerical experiments are conducted to show that both OCBA+ and OCBAR can achieve higher PCS than the original OCBA under the same budget.

2 CONVERGENCE RATE ANALYSIS

Following the standard assumption that the simulation outputs are independent normal samples, we investigate the performance of algorithms from a large deviations (LD) perspective. Due to the complexity of analysis, our results mainly focus on a two-design case. Two algorithms, OCBA+ and OCBAR, are proposed as variants of OCBA. They both allocate the simulation budget in a fully sequential style, where at each iteration, OCBA+ simulates the design that has the largest gap between target allocation and actual allocation, and OCBAR decides which design to simulate by sampling from a probability distribution constructed from target allocations. Our major discovery is that if both algorithms use a constant initial sample size (independent of the total budget T) for estimating the designs mean and variance, then the PFS (defined as 1-PCS) will converge to 0 only at a sub-exponential rate. More specifically, we have the following result.

Theorem 1 Under some regularity assumptions, if OCBA+ and OCBAR use an initial sample size that is independent of the total budget , then in the case of two designs,

1. For OCBA+, the PFS is lower bounded by a polynomial function of T.
2. For OCBAR, the LD rate of the PFS is 0.

Theorem 1 is somewhat surprising, because it implies that (i) the common practice of a constant initial sample size suffers from a slow convergence rate; (ii) even if the allocation fractions converge to the optimal ones as , the LD rate of the PFS could be drastically different from the optimal rate. An intuitive explanation is that during the initialization phase, the designs mean and variance estimates are subject to some rare but extreme errors, which will cause the algorithms to no longer allocate budget to some designs (also called “freezing”) and eventually result in a false selection. It can be shown that such an event is not exponentially rare. The same proof technique can be applied to show a similar result for the original OCBA, and it is promising to extend the results beyond two designs.

3 IMPROVING OCBA

In light of Theorem 1, we modify OCBA and its variants by forcing the initial sample size to grow linearly in T. This modification allows us to derive a finite-sample bound on the PFS, which guarantees an exponential convergence rate. Furthermore, we use numerical experiments to demonstrate the improvement achieved by the modified algorithms, where the PFS is shown to be smaller than that of their original counterparts under every finite T.

REFERENCES

Chen, C.-H., J. Lin, E. Yücesan, and S. E. Chick. 2000. “Simulation Budget Allocation for Further Enhancing the Efficiency of Ordinal Optimization”. *Discrete Event Dynamic Systems* 10(3):251–270.