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

FULLY SEQUENTIAL RANKING AND SELECTION PROCEDURES WITH PAC GUARANTEE

Ying Zhong
L. Jeff Hong

Department of Management Sciences, City University of Hong Kong
Kowloon, HONG KONG

ABSTRACT

In the ranking-and-selection field, many current existing fully sequential procedures are developed under the indifference-zone (IZ) formulation which assumes an optimality gap between the best alternative and the others. In this paper, by modifying one classical fully sequential procedure, Paulson's procedure, we devise a new type of fully sequential procedure which can provide the probably approximately correct (PAC) selection guarantee. The procedure is able to select an alternative within a certain distance from the best one with the desired probability even when the IZ assumption is violated. In order to provide this selection guarantee, we lose some efficiency on the procedure. However, we show, both theoretically and numerically, that this efficiency loss is limited and acceptable when the number of alternatives is large.

1 INTRODUCTION

In the simulation literature, using simulation experiments to select the alternative with the best mean performance among a finite group of candidates is often referred to as the ranking-and-selection (R&S) problem, see Bechhofer et al. (1995), CKim and Nelson (2006), and Chick (2006) for reviews. Based on the statistical inference being used, the current existing procedures can be classified into two approaches: the Bayesian approach and the frequentist approach. The objectives of these two classes of procedures are quite different. Bayesian procedures, e.g.: Chick and Inoue (2001) and Chick and Frazier (2012), aiming to minimize the expected opportunity cost, typically have high efficiencies to solve practical problems, but generally, do not provide any frequentist's statistical guarantees. Meanwhile, for frequentist procedures, e.g.: Rinott (1978), Paulson (1964), Kim and Nelson (2001), and Fan et al. (2016), offering certain kind of statistical guarantee is always the primary goal. Many frequentist procedures are designed to deliver the desired *probability of correct selection* (PCS), even under the least favorable configuration of the means. Therefore, comparing with Bayesian procedures, frequentist procedures are averagely conservative. In this work, we mainly focus on developing frequentist procedures.

Under the frequentist framework, the goal of selecting the unique best is stringent, because the mean performance of second best alternative can be arbitrarily close to that of the best one. To avoid such difficulty, many frequentist procedures adopt the *indifference-zone* (IZ) formulation, which is initially proposed by Bechhofer (1954). It assumes that the mean difference between the best and the second best is at least $\delta > 0$, where δ is the smallest difference the decision makers feel worth detecting. By suggesting an optimality gap between the best alternative and the others, the IZ assumption ensures that the best alternative is identifiable in a finite number of samples.

So far, under the IZ formulation, various types of procedures have been proposed, ranging from the earliest two-stage procedures, e.g.: Rinott (1978), which are designed to handle problems with $k \leq 100$, to the latter fully sequential procedures, which are capable of solving medium size problems (e.g., $k \leq 500$) by eliminating clearly inferior alternatives in the middle of the selection process. Recently, with the continuing development of the field, there are growing needs for the possibilities to solve even large-scale R&S problems (i.e., $k \geq 10^4$). As pointed out by Luo et al. (2015), the R&S problem can fit easily into the

parallel computing environments. Both Luo et al. (2015) and Ni et al. (2017) make the attempts to solve large-scale R&S problems by implementing some variants of current existing procedures in the parallel computing environments.

A natural question, then, arises: whether the IZ assumption remains solid as the number of alternatives gets larger and larger? The answer is negative. As illustrated by Ni et al. (2017), when the number of alternatives is large, there always can exist many δ-optimal alternatives which lie within δ of the best alternative. To address the issue, Ni et al. (2017) propose a new form of selection guarantee which requires the procedures to select a δ-optimal alternative with the predefined probability $1-\alpha$ even when the IZ assumption is violated. Ma and Henderson (2017) call it the *probably approximately correct* (PAC) guarantee. It turns out that, even though the fully sequential procedures inherently are the best candidates procedures for solving problems with large k, there is no rigorous proof by far stating that fully sequential procedures satisfy the PAC guarantee. Without the IZ assumption, the statistical validity of these fully sequential procedures are questionable, contradicting the primary goal of frequentist procedures.

In this paper, we tackle this problem by redesigning one classical fully sequential procedure, Paulson's procedure. We show that after the modifications, the new procedure not only maintains the fully sequential characteristic but also satisfies the PAC guarantee. To achieve this, we want to notice that, the failure of proving the PAC guarantee for traditional fully sequential procedures is largely due to the fact that, in the presence of δ-optimal alternatives, the probability that best alternative is eliminated by another δ-optimal alternative in the middle of the selection process is not well controlled and is relatively high. After the best alternative is eliminated, the procedures run without any statistical guarantee. In light of this, we properly widen the continuation regions, which are used to conduct pairwise comparisons between alternatives, in Paulson's procedure to grant the best alternative more chance of surviving in the competitions with the other δ-optimal alternatives. Essentially, we satisfy the PAC guarantee at the expense of the procedure's efficiency loss. Nevertheless, we will show, both theoretically and numerically, that this efficiency loss is limited and acceptable when the number of alternatives is large.

The rest of the paper is organized as follows. In section 2, with a simple setting, where all the alternatives share a common known variance, we demonstrate our main idea of designing the procedure and conduct the efficiency analysis on the procedure. In section 3, we extend our procedure to the case of unknown and unequal variances followed by some preliminary numerical results in section 4.

2 PROCEDURE FOR A COMMON KNOWN VARIANCE

To facilitate our analysis, we first introduce some standard notations. Let $\mathbb{K} = \{1,2,\ldots,k\}$ be the set of alternatives. For each alternative $i \in \mathbb{K}$, it generates independent and identically distributed (i.i.d.) observations $X_{i,\ell}$ for $\ell = 1,2,\ldots$, from the normal distribution with mean μ_i and variance σ_i^2. $\bar{X}_i(t)$ is defined as the sample average of alternative i based on t observations. Without loss of generosity, we assume that the means are in an ascending order, i.e., $\mu_1 \leq \mu_2 \leq \cdots \leq \mu_k$, and alternative k is the best alternative. Being different from the IZ formulation, we make no assumption on the configuration of the means. The parameter $\delta > 0$ is defined as the PAC parameter in the rest of the paper. We call an alternative i δ-optimal, if $\mu_k - \mu_i \leq \delta$. It is possible that whenever the user specifies a positive constant $\delta > 0$, there is more than one δ-optimal alternative. Similarly, an alternative j is called δ-away, if $\mu_k - \mu_j > \delta$. Our objective is to devise a procedure that can achieve,

$$\mathbb{P}(\mu_k - \mu_I \leq \delta) \geq 1-\alpha, \tag{1}$$

where I is index of the alternative the procedure terminates with, α is the *probability of incorrect selection* (PICS), and $1-\alpha$ is called the PAC probability in this paper. It means that the procedure needs to produce a δ-optimal alternative with the confidence greater or equal to $1-\alpha$. We call the procedure satisfying equation (1) the (δ,α)-PAC procedure. If the IZ assumption indeed holds, i.e., $\mu_1 \leq \mu_2 \leq \cdots < \mu_k - \delta$, for a (δ,α)-PAC procedure, we have,

$$\mathbb{P}(\mu_k - \mu_I \leq \delta | \mu_k > \mu_{k-1} + \delta) = \mathbb{P}(I = k | \mu_k > \mu_{k-1} + \delta) \geq 1-\alpha,$$

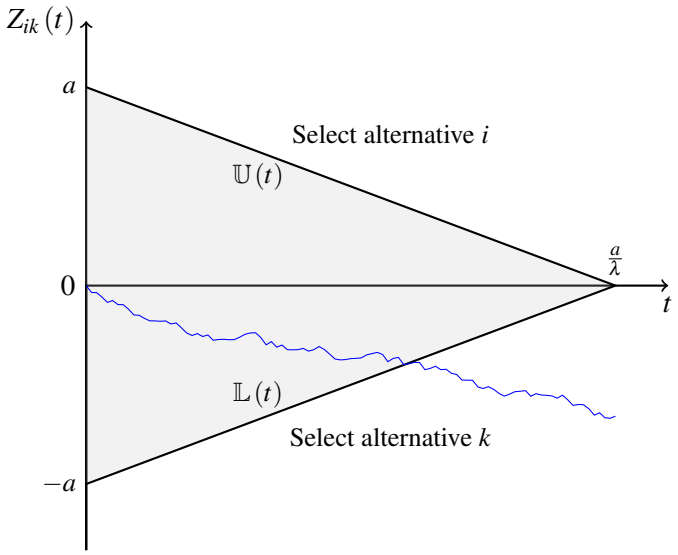

Figure 1: Continuation Region for Paulson's Procedure.

indicating that the (δ,α)-PAC procedure can still guarantee to select the unique best alternative k, and the PAC probability is equivalent to the classical PCS under this situation. In the following subsections, with a brief review on Paulson's procedure, we show how Paulson's procedure can be modified to satisfy the PAC guarantee.

2.1 Paulson's Procedure

For the traditional fully sequential procedures, it is a convention to break down the problem of selecting the best alternative to subproblems, which make pairwise comparisons between every two competing alternatives. The comparison between two alternatives i and j refers to their partial sum difference process, $\left\{Z_{ij}(t) = t\cdot(\bar{X}_i(t)-\bar{X}_j(t)) : t = 1,2,\ldots\right\}$. Once the random process $Z_{ij}(t)$ exists its predetermined continuation region, a decision can be made. Usually, the continuation regions are constructed to ensure that, for each comparison between the best alternative (i.e., alternative k) and the other $k-1$ alternatives, falsely eliminating alternative k happens with probability $\alpha/(k-1)$, so that,

$$\mathbb{P}(k \text{ is eliminated}) = \sum_{i=1}^{k-1}\mathbb{P}(i \text{ eliminates } k) \le (k-1)\times\frac{\alpha}{k-1} = \alpha.$$

There are various ways of constructing the continuation regions. Under the normality assumption, one common way is to model the partial sum difference process as a Brownian motion (BM) process, and use the theories of hitting times in the BM process to approximate these continuation regions.

We take Paulson's procedure as an example to show the underlying mechanisms of constructing these continuation regions. Particularly, we consider the comparison between two alternatives i and k, where $i \neq k$. As shown by Figure 1, the selection of these two alternatives depends on whether $Z_{ik}(t)$ exits the triangular continuation region from the lower decision boundary $\mathbb{L}(t) = -a+\lambda\cdot t$ or the upper decision boundary $\mathbb{U}(t) = a-\lambda\cdot t$. The false elimination occurs if $Z_{ik}(t)$ exists the continuation region from $\mathbb{U}(t)$. From the BM process point of view, the theoretical foundations of constructing the continuation region (i.e., determining the decision boundaries) are twofold: 1) Throughout this paper, we let $\mathscr{B}(\cdot)$ denote the BM process with unit variance, and $\mathscr{B}_\Delta(\cdot)$ is defined as the BM process with unit variance and drift Δ. As suggested by Theorem 1 of Hong (2006), for any user inputted constant $\lambda \in (0,\delta)$, the random processes $\left\{(Z_{ik}(t)+\lambda\cdot t)/\sigma_{ik}^2 : t = 1,2,\ldots\right\}$ and $\left\{\mathscr{B}_{\mu_i-\mu_k+\lambda}\left(t/\sigma_{ik}^2\right) : t = 1,2,\ldots\right\}$ have the same joint distribution,

where $\sigma_{ik}^2 = Var(X_{i,\ell} - X_{k,\ell})$; 2) If the IZ assumption holds, $\mathscr{B}_{\mu_i-\mu_k+\lambda}\left(t/\sigma_{ik}^2\right)$ always has a drift less or equal to $\lambda - \delta$. A well known result for the BM process with a negative drift bounded away from zero is summarized in Lemma 1.

Lemma 1 For a BM with unit variance and drift $\Delta < 0$, if $a' > 0$,

$$\mathbb{P}\left(\sup_{t\geq 0}\mathscr{B}_{\Delta}(t) \geq a'\right) = \exp\left(-2\left|\Delta\right|\left|a'\right|\right).$$

Therefore, if σ_{ik}^2 is known, by simply letting $a = \sigma_{ik}^2 \cdot \log\left[(k-1)/\alpha\right]/\left[2\cdot(\delta-\lambda)\right]$, the upper decision boundary $\mathbb{U}(t)$ with the form $a - \lambda \cdot t$ satisfies,

$$\mathbb{P}(i \text{ eliminates } k) = \mathbb{P}\left(\sup_{t\geq 1} Z_{ik}(t) \geq \mathbb{U}(t)\right) \leq \frac{\alpha}{k-1}.$$

In the meantime, the lower decision boundary $\mathbb{L}(t)$ is, by symmetry of the random processes $Z_{ik}(t)$ and $Z_{ki}(t)$, the reflection of $\mathbb{U}(t)$ over x-axis.

The IZ assumption is crucial. Without the IZ assumption, the BM process $\mathscr{B}_{\mu_i-\mu_k+\lambda}\left(t/\sigma_{ik}^2\right)$ may have a drift arbitrarily close to zero or even a positive drift. Under this situation, with high probability, alternative k could be eliminated by alternative i during the selection process. Moreover, with the equally high probability, alternative i could be further eliminated by another alternative j which is δ away from alternative k but within δ of alternative i, leading the procedure to finally choose a δ-away alternative. In view of this, while constructing the continuation region for our (δ,α)-PAC procedure, we add an additional negative drift to the random process $Z_{ik}(t) + \lambda \cdot t$. Particularly, we consider the random process $Z_{ik}(t) + (\lambda - \delta)\cdot t$ and its corresponding BM process $\mathscr{B}_{\mu_i-\mu_k+\lambda-\delta}\left(t/\sigma_{ik}^2\right)$. Thus, the new process always has a negative drift bounded away from zero. As we will show in the following subsection, this small change leads to a wider continuation region. Consequently, the probability that alternative k is eliminated by alternative i can be well controlled regardless whether alternative i is δ-optimal or not.

2.2 Procedure 1

In this subsection, we introduce a fully sequential (δ,α)-PAC procedure under the setting of a common known variance. This is in line with the setting when Paulson's procedure was first proposed.

Assumption 1 (A Common Known Variance) For each alternative $i \in \mathbb{K}$, it generates independently and identically distributed (i.i.d.) observations $X_{i,\ell}$, for $\ell = 1,2,\ldots$, from the normal distribution with unknown mean μ_i and known variance σ^2. The observations generated by different alternatives are independent.

Our procedure operates in almost the same logic as the classical Paulson's procedure, except the continuation regions the procedure adopts to conduct pairwise comparisons. Firstly, the procedure constructs the continuation regions for all the possible pairwise comparisons based on the common known variance. Then, in the elimination stage, at each round, the procedure screens whether there is a partial sum difference process $Z_{ij}(t)$ exiting its continuation region. If it is the case, the procedure eliminates one alternative accordingly. After all the eliminations, the procedure collects one more observation for each surviving alternative and proceeds to the next round of screening. The procedure stops when there is only one alternative left or the number of observations for each surviving alternative reaches the maximum. The detailed descriptions of the procedure are listed below.

Procedure 1 (Normal Observations with a Common Known Variance)

Input: Select the PAC probability $1-\alpha$ $(0 < \alpha \leq 1 - 1/k)$. Set parameters $\delta > 0$ and λ $(0 < \lambda < \delta)$.

Initialization: Let $\mathbb{I} = \{1,2,\ldots,k\}$ be the set of alternatives in contention. For each $i \in \mathbb{I}$, simulate one observation $X_{i,1}$. Compute

$$a = \frac{\sigma^2 \log\left[(k-1)/\alpha\right]}{(\delta-\lambda)} \text{ and } N = \left\lfloor \frac{a}{\lambda} \right\rfloor + 1,$$

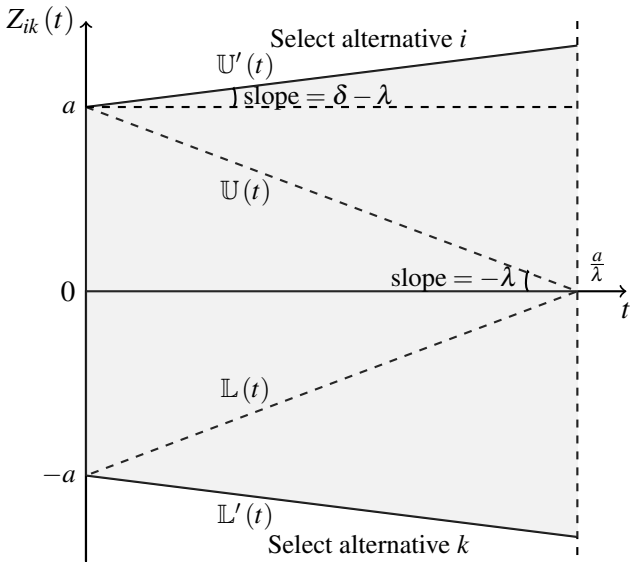

Figure 2: Continuation Region for Procedure 1.

where $\lfloor \cdot \rfloor$ indicate truncation of any fractional part, and set $t = 1$.
Screening: Set $\mathbb{I}^{old} = \mathbb{I}$. Let

$$\mathbb{I} = \left\{ i : i \in \mathbb{I}^{old} \text{ and } t \cdot (\bar{X}_i(t) - \bar{X}_j(t)) \geq -a - (\delta - \lambda) \cdot t, \forall j \in \mathbb{I}^{old}, j \neq i \right\}.$$

Stopping Rule:

- If $|\mathbb{I}| = 1$, stop and select the alternative whose index is in $\mathbb{I}$ as the best.
- Else if $t = N$, stop and select the alternative whose index is in $\mathbb{I}$ having the largest $\bar{X}_i(t)$ as the best.
- Otherwise, take one additional observation $X_{i,t+1}$ from each alternative $i \in \mathbb{I}$, set $t = t + 1$, and go to **Screening**.

Remark 1 Recall that, for Paulson's procedure, **Screening** conducts eliminations with milder conditions, i.e.,

$$\mathbb{I} = \left\{ i : i \in \mathbb{I}^{old} \text{ and } t \cdot (\bar{X}_i(t) - \bar{X}_j(t)) \geq -a + \lambda \cdot t, \forall j \in \mathbb{I}^{old}, j \neq i \right\}.$$

To better illustrate the changes we have made on Paulson's procedure, in Figure 2, we plot the decision boundaries $\mathbb{U}'(t) = a + (\delta - \lambda) \cdot t$ and $\mathbb{L}'(t) = -a - (\delta - \lambda) \cdot t$ for the comparison between alternatives i and k in Procedure 1. It can be observed that, there are two distinct features for the continuation region in Procedure 1: 1) As we change the random process used to construct the continuation region, the decision boundaries of Paulson's procedure have opened up; 2) For Paulson's procedure, the continuation region naturally vanishes at a/λ, where $\mathbb{U}(t)$ and $\mathbb{L}(t)$ intersect. As in Procedure 1, there is no such an intersection point, we force the continuation region in Procedure 1 to end at the same point as Paulson's procedure does.

These modifications are quite intuitive. On the one hand, after we widen all the continuation regions, the new regions are uniformly larger than those of Paulson's procedure. It suggests that the procedure becomes more conservative, and requires more observations to eliminate alternatives. This is true, because the procedure imposes tighter conditions to eliminate alternatives, lowering the chance that the best one

is falsely eliminated. On the other hand, forcing all the pairwise comparisons to stop at a/λ ensures that every δ-away alternative can be properly eliminated. From these two aspects, we are able to establish the statistical validity for Procedure 1.

Theorem 1 If Assumption 1 holds, Procedure 1 is a (δ,α)-PAC procedure.

2.3 Asymptotic Efficiency

Besides the effectiveness (statistical validity), the efficiency is also an important aspect to measure the performance of a procedure. In this subsection, we establish the efficiency analysis for Procedure 1. Specifically, we consider the efficiency loss of Procedure 1, compared with Paulson's procedure, while solving a problem where the IZ assumption holds. As Procedure 1 is applied, we exam the expected sample size alternative i can take for $i = 1,2,\ldots,k-1$, in the asymptotic regime of letting α go to zero, and compare the expected sample size with that in Paulson's procedure. The main technique we use is the mean path analysis provided by Perng (1969).

Theorem 2 Suppose that there are in total k alternatives. For each alternative $i \in \mathbb{K}$, it generates i.i.d. observations $X_{i,\ell}$, for $\ell = 1,2,\ldots$, from the normal distribution with unknown mean μ_i and known variance σ^2. The observations generated by different alternatives are independent. Let $\mathbb{N}_i$ be the sample size alternative i can take in Paulson's procedure, i.e.,

$$\mathbb{N}_i = \min\left\{t_j : \min_{j\in\mathbb{I}_{t-1}} Z_{ij}(t) \le -a - \delta\cdot t\right\},$$

and $\mathbb{N}_i'$ be the expected sample size alternative i can take in Procedure 1, i.e.,

$$\mathbb{N}_i' = \min\left\{t : \min_{j\in\mathbb{I}_{t-1}} Z_{ij}(t) \le -a - (\delta-\lambda)\cdot t \text{ or } t = N\right\},$$

where $\mathbb{I}_{t-1}$ is the set of alternatives that have not been eliminated up to time $t-1$. If $\mu_1 \le \mu_2 \le \cdots \le \mu_{k-1} < \mu_k - \delta$, $1-\alpha$ is defined as the PAC probability, and both procedures are applied, then,

$$\lim_{\alpha\to 0}\frac{\mathbb{E}[\mathbb{N}_i']}{\mathbb{E}[\mathbb{N}_i]} = \frac{\mu_k-\mu_i+\lambda}{\mu_k-\mu_i-\delta+\lambda}. \tag{2}$$

It can be seen that, for equation (2), the ratio on the right hand side quantifies the efficiency loss of Procedure 1. The larger the ratio the larger loss on the efficiency. As the ratio is always greater than one, it means that there is always an efficiency loss on Procedure 1. This efficiency loss can be thought as the price that we pay for the PAC guarantee. For any fixed λ, when $\mu_k-\mu_i=\delta$, the largest efficiency loss is incurred. Under this situation, the sample size alternative i can take in Procedure 1 is $(\lambda+\delta)/\lambda$ times as that in Paulson's procedure. However, when the mean difference $\mu_k-\mu_i$ gets large, the ratio quickly approaches to one, implying that the difference in efficiency shrinks. It is arguable that when the number of alternatives is large, most of the alternatives are clearly inferior ones, and the sample size needed to eliminate these clearly inferior alternatives are almost the same for both procedures. Thus, the total sample size of Procedure 1 would be similar to that of Paulson's procedure while solving a large-scale R&S problem.

3 PROCEDURE FOR UNKNOWN AND UNEQUAL VARIANCES

As the setting of Procedure 1 can hardly be achieved in practice, in this section, we extend Procedure 1 to the case of unknown and unequal variances.

3.1 Procedure 2

In most simulation studies, variances of the outputs are typically unknown and unequal, and need to be estimated during the selection process. In this subsection, we introduce a procedure, namely, Procedure 2, that relaxes the assumption of a common known variance in Procedure 1.

Assumption 2 (Unknown and Unequal Variances) For each alternative $i \in \mathbb{K}$, it generates independently and identically distributed (i.i.d.) observations $X_{i,\ell}$, for $\ell = 1,2,\ldots$, from the normal distribution with unknown mean μ_i and unknown variance σ_i^2.

Remark 2 In Assumption 2, we do not emphasize the independence between the observations generated by different alternatives. It allows the use of the common random numbers (CRNs).

Compared with Procedure 1, Procedure 2 additionally requires the first stage sampling to estimate the unknown variances for all the alternatives. Once the variance information is collected, the procedure uses the first stage sample variances to construct the continuation regions needed in the second stage. Then, the eliminations are conducted in a similar manner to that of Procedure 1. Note that, as the variances of the alternatives are unequal, the continuation regions in Procedure 2 usually differ from one pairwise comparison to another. Thus, how to determine the ending point of each continuation region becomes a problem. In terms of the procedure's efficiency, it would be appealing to force different continuation regions to end at different points, making full use of the variance information. However, it directly leads the procedure to the same problem as that encountered by the traditional fully sequential procedures while satisfying the PAC guarantee. The probability that alternative k is eliminated by another δ-optimal alternative in the middle of the selection process cannot be controlled. Therefore, in Procedure 2, we keep adopting a universal ending point for all the continuation regions. We force all the continuation regions to end at the point where the worst-case scenario does. The details of the procedure are listed below.

Procedure 2 (Normal Observations with Unknown and Unequal Variances)

Input: Select the PAC probability $1-\alpha$ $(0 < \alpha \leq 1-1/k)$. Set parameters $\delta > 0$ and λ $(0 < \lambda < \delta)$.

Initialization: Let $\mathbb{I} = \{1,2,\ldots,k\}$ be the set of alternatives in contention. For each $i \in \mathbb{I}$, simulate n_0 observations $X_{i,1}, X_{i,2}, \ldots, X_{i,n_0}$ from alternative i. Let $h^2 = \frac{n_0-1}{4(\delta-\lambda)} \times \left[\left(\frac{\alpha}{k-1}\right)^{-\frac{2}{n_0-1}} - 1\right]$. For all $i \neq j$, compute

$$S_{ij}^2 = \frac{1}{n_0-1}\sum_{\ell=1}^{n_0}\left[X_{i,\ell} - X_{j,\ell} - (\bar{X}_i(n_0) - \bar{X}_j(n_0))\right]^2 \text{ and } N_{ij} = \left\lfloor \frac{h^2 S_{ij}^2}{\lambda} \right\rfloor + 1.$$

Let

$$N_i = \max_{j \neq i} N_{ij}.$$

Set the $t = n_0$.

Screening: Set $\mathbb{I}^{old} = \mathbb{I}$. Let

$$\mathbb{I} = \left\{ i : i \in \mathbb{I}^{old} \text{ and } t \cdot (\bar{X}_i(t) - \bar{X}_j(t)) \geq -h^2 S_{ij}^2 - (\delta-\lambda)\cdot t, \forall j \in \mathbb{I}^{old}, j \neq i \right\}. \tag{3}$$

Stopping Rule:

- If $|\mathbb{I}| = 1$, stop and select the alternative whose index is in $\mathbb{I}$ as the best.
- Else if $t = \max_{i \in \mathbb{K}} N_i$, stop and select the alternative whose index is in $\mathbb{I}$ having the largest $\bar{X}_i(t)$ as the best.
- Otherwise, take one additional observation $X_{i,t+1}$ from each alternative $i \in \mathbb{I}$, set $t = t+1$, and go to **Screening**.

Notice that one may argue that if the variances are unequal, our procedure may experience the conservativeness induced by forcing all the continuation regions to end at the point where the worst-case scenario does. The technique proposed by Hong (2006) can properly address this issue. By adjusting the sampling rule, we are not necessary to sample every surviving alternatives the same amount of observations at each round. We are able to allocate different numbers of observations to different alternatives based on their variances (or, first stage sample variances) so as to rescale all the continuation region to about the same size. Like many fully sequential procedures, updating the variance information for each alternative in the second stage is generally not allowed in Procedure 2. The procedure requires the independence between the sample variances and the overall sample means in the second stage to guarantee the finite time statistical validity. In the theorem below, we summarize the statistical validity for Procedure 2.

Theorem 3 If Assumption 2 holds, Procedure 2 is a (δ, α)-PAC procedure.

Except the designing of the ending point of each continuation region, the other techniques we used to extend Procedure 1 to handle the case of unknown and unequal variances are quite standard. Interested readers may also see Paulson (1964) and Hong (2006) for more details on how these techniques.

3.2 Speed Up Pairwise Comparisons

Many current existing R&S procedures are originally designed for the purpose of solving medium size problems (i.e., $k \leq 500$). After years of rapid growth of the computing power, many recent works start to investigate the possibilities of employing these procedures to solve large-scale R&S problems in the parallel computing environments. For instance, Luo et al. (2015) attempt to use KN procedure proposed by Kim and Nelson (2001) to solve the problem with more than 10^4 alternatives, and Ni et al. (2017) even push the number up to 10^6. Both studies reveal the fact that when the number of alternatives gets large, the nature of all pairwise comparisons at each round of screening becomes a bottleneck in a fully sequential procedure. The all pairwise comparisons at each round generally require a computational complexity of $\mathcal{O}\left(k^2\right)$ which is significantly larger than the computational complexity of $\mathcal{O}(k)$ needed to simulate new observations in each round of the comparisons.

Without considering the CRNs, Hong et al. (2016) propose a simple but effective technique to reduce the computational complexity for the all pairwise comparisons at each round to the order of $\mathcal{O}(k)$ for some fully sequential procedures. This technique is applicable to Procedure 2 as well. The general idea works as follows. With the independence between observations generated by different alternatives, we can replace S_{ij}^2 in Procedure 2 by the term $S_i^2 + S_j^2$ where $S_i^2 = \frac{1}{n_0-1}\sum_{\ell=1}^{n_0}\left(X_{i,\ell} - \bar{X}_i(n_0)\right)^2, \forall i \in \mathbb{K}$. If we further let

$$W_i^+(t) = t \cdot \bar{X}_i(t) + h^2 S_i^2 \text{ and } W_i^-(t) = t \cdot \bar{X}_i(t) - h^2 S_i^2, \forall i \in \mathbb{I},$$

and $i_-^* = \arg\max_{i \in \mathbb{I}} W_i^-(t)$ at time t, then the screening process (3) in Procedure 2 can be rephrased to

$$\mathbb{I} = \left\{ i : i \in \mathbb{I}^{old} \text{ and } W_i^+(t) \geq W_{i_-^*}^-(t) - (\delta - \lambda) \cdot t \right\}. \tag{4}$$

It suggests that, at each round of screening, one does not need to conduct all pairwise comparisons among the surviving alternatives; instead one can define a "best" alternative based on certain criteria, and only compare each surviving alternative with the "best" alternative. Thus, the order of comparisons at each round reduce to $\mathcal{O}(k)$. The proposition below states that after such modifications, the statistical validity of Procedure 2 remains intact. Proposition 1 provides a theoretical guarantee to apply our procedure to solve large-scale R&S problems.

Proposition 1 If Assumption 2 holds, without considering the CRNs, as we replace the term S_{ij}^2 by $S_i^2 + S_j^2$, and the screening process (3) by (4), Procedure 2 remains a (δ, α)-PAC procedure.

4 NUMERICAL RESULTS

For the numerical experiments listed below, by default, the desired PAC probability, $1-\alpha$, is set to be 0.95, and we let $\delta = 0.1$. We first consider a simple example where there are two alternatives sharing a common known variance, i.e., $\sigma_1^2 = \sigma_2^2 = \sigma^2 = 1$. We compare the performances of Procedure 1, Paulson's procedure, and another famous fully sequential procedure, KN procedure, under different configurations of the means. We summarize the estimated PAC probability and the average total sample size with 95% confidence interval over 1,000 macro-replications in Table 1.

Table 1: Comparisons of Procedure 1, Paulson's Procedure, and KN Procedure.

$\mu_2-\mu_1$	Procedure 1: λ		Paulson's procedure: λ		KN: c
	$\delta/3$	$\delta/2$	$\delta/3$	$\delta/2$	1
$\mu_2-\mu_1=\delta$	1.00 1.70×10^3 $\pm0.05\times10^3$	0.99 1.72×10^3 $\pm0.04\times10^3$	0.97 0.66×10^3 $\pm0.02\times10^3$	0.98 0.79×10^3 $\pm0.02\times10^3$	0.96 0.60×10^3 $\pm0.02\times10^3$
$\mu_2-\mu_1=2\delta$	1.00 0.66×10^3 $\pm0.02\times10^3$	1.00 0.79×10^3 $\pm0.02\times10^3$	1.00 0.38×10^3 $\pm0.01\times10^3$	1.00 0.49×10^3 $\pm0.01\times10^3$	1.00 0.38×10^3 $\pm0.01\times10^3$
$\mu_2-\mu_1=4\delta$	1.00 0.27×10^3 $\pm0.01\times10^3$	1.00 0.34×10^3 $\pm0.01\times10^3$	1.00 0.21×10^3 $\pm0.01\times10^3$	1.00 0.27×10^3 $\pm0.01\times10^3$	1.00 0.21×10^3 $\pm0.01\times10^3$
$\mu_2-\mu_1=8\delta$	1.00 0.13×10^3 $\pm0.01\times10^3$	1.00 0.16×10^3 $\pm0.01\times10^3$	1.00 0.11×10^3 $\pm0.01\times10^3$	1.00 0.14×10^3 $\pm0.01\times10^3$	1.00 0.11×10^3 $\pm0.01\times10^3$

For Procedure 1 and Paulson's procedure, the smaller λ we set, the smaller total sample sizes are required to identify the best alternative (see, each row). This occurs because, for these two procedures, a small λ usually associates with tight continuation regions when t is small. As a result, the partial sum difference processes may stop early in expectation. Meanwhile, this is at the expense of larger ending points for the continuation regions, i.e., a worse performance in the worst-case scenario. It also can be seen that, compared with the other two traditional fully sequential procedures, Procedure 1 requires larger total sample sizes to solve the problems. This is because KN procedure in general provides the tightest continuation regions among the traditional fully sequential procedures by employing the Fabian (1974)'s result, and the continuation regions in Procedure 1 are always larger than those in Paulson's procedure. This is the efficiency loss incurred by satisfying the PAC guarantee. However, as the difference in means, i.e., $\mu_2-\mu_1$, increases, the difference in efficiency diminishes. When $\mu_2-\mu_1=8\delta$ (the last row), the total sample sizes needed to solve the problem are almost the same for these three procedures. As in practice, for large-scale R&S problems, the means of the alternatives are often spread out over a wide range. Most of the alternatives are clearly inferior ones. It suggests that while solving a large-scale R&S problem, the performances of these three procedures would not differ far away.

We then exam the performances of these three procedures with a large-scale R&S problem. Particularly, we consider a problem with 10^4 alternatives. In this example, for each replication, the means of the alternatives are drawn from the normal distribution with mean 0 and variance $16\delta^2$, i.e., $\mu_i \sim N\left(0,16\delta^2\right)$ for $i=1,2,\ldots,10^4$, and are i.i.d. The variances of the alternatives are equal, $\sigma_1^2=\sigma_2^2=\cdots=1$, but unknown. Thus, the unknown-variance versions of these three procedures are applied. We also let different alternatives generate independent observations, enabling the use of the technique provided in section 3.2 to reduce the computational complexity of comparisons for all the three procedures. We estimate the interested statistics based on 100 macro-replications. In Table 2, we report the average number of δ-

optimal alternatives appeared in each replication. For the performance measure, we report estimated the PAC probability, and the average total sample size for each procedure.

Table 2: Comparisons of Procedure 2, Paulson's Procedure, and KN Procedure.

k	# of δ-optimal alternatives	Procedure 2: $\lambda = \delta/2$	Paulson's Procedure: $\lambda = \delta/2$	KN: $c = 1$
10^4	2.98	1.00 3.02×10^6	1.00 2.92×10^6	1.00 2.71×10^6

From Table 2, we have several findings. First, on average, for each replication, there exist around 3 δ-optimal alternatives. It suggests that the IZ assumption is no longer appropriate for a large-scale R&S problem. Second, all the three procedures over-deliver the PAC probability in our experiment. This is due to the fact that, the continuation regions provided by Procedure 2 actually provides a theoretical lower bound to the PAC guarantee. Because the frequentist procedures are always designed in the most conservative ways, for most of the cases, the traditional fully sequential procedures can numerically demonstrate the robustness of PAC guarantee. However, without the theoretical guarantee, the statistical validity of these procedures is unclear. Third, there are about 3.4% and 11.4% inflation on the total sample size for Procedure 2 compared with Paulson's procedure and KN procedure respectively, and we think it is tolerable. This is consistent with our previous analysis that the efficiency of these three procedures are similar while solving a large-scale R&S problem.

5 CONCLUSION

In this paper, we provide a new fully sequential procedure that can satisfy the PAC guarantee by proper widening the continuation regions of the classical Paulson's procedure. We show, both theoretically and numerically, that the new procedure is suitable for solving large-scale R&S problems.

REFERENCES

Bechhofer, R. E. 1954. "A Single-sample Multiple Decision Procedure for Ranking Means of Normal Populations with Known Variances". *The Annals of Mathematical Statistics* 25(1):16–39.

Bechhofer, R. E., D. M. Goldsman, and T. J. Santner. 1995. *Design and Analysis of Experiment for Statistical Selection, Screening, and Multiple Comparisons*. 1st ed. New York: Wiley.

Chick, S. E., and K. Inoue. 2001. "New Two-stage and Sequential Procedures for Selecting the Best Simulated System". *Operations Research* 49(5):732–743.

Chick, S. E. 2006. "Subjective Probability and Bayesian Methodology". In *Handbooks in Operations Research and Management Science: Simulation*, edited by S. G. Henrson and B. L. Nelson, Chapter 9, 225–257. Oxford, UK: Elsevier Science.

Chick, S. E., and P. Frazier. 2012. "Sequential Sampling with Economics of Selection Procedures". *Management Science* 58(3):550–569.

Fabian, V. 1974. "Note on Anderson's Sequential Procedures with Triangular Boundary". *The Annals of Statistics* 2(1):170–176.

Fan, W., L. J. Hong, and B. L. Nelson. 2016. "Indifference-zone-free Selection of the Best". *Operations Research* 64(5):1499–1514.

Hong, L. 2006. "Fully Sequential Indifference-zone Selection Procedures with Variance-dependent Sampling". *Naval Research Logistics (NRL)* 53(5):464–476.

Hong, L. J., J. Luo, and Y. Zhong. 2016. "Speeding Up Pairwise Comparisons for Large Scale Ranking and Selection". In *Proceedings of the 2016 Winter Simulation Conference*, edited by T. M. K. Roeder et al., 749–757. Piscataway, New Jersey: IEEE.

Kim, S.-H., and B. L. Nelson. 2001. "A Fully Sequential Procedure for Indifference-zone Selection in Simulation". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 11(3):251–273.

Kim, S.-H., and B. L. Nelson. 2006. "Selecting the Best System". In *Handbooks in Operations Research and Management Science: Simulation*, edited by S. G. Henrson and B. L. Nelson, Chapter 17, 501–534. Oxford, UK: Elsevier Science.

Luo, J., L. J. Hong, B. L. Nelson, and Y. Wu. 2015. "Fully Sequential Procedures for Large-scale Ranking-and-selection Problems in Parallel Computing Environments". *Operations Research* 63(5):1177–1194.

Ma, S., and S. G. Henderson. 2017. "An Efficient Fully Sequential Selection Procedure Guaranteeing Probably Approximately Correct Selection". In *Proceedings of the 2017 Winter Simulation Conference*, edited by W. K. V. Chan et al., 2225–2236. Piscataway, New Jersey: IEEE.

Ni, E. C., D. F. Ciocan, S. G. Henderson, and S. R. Hunter. 2017. "Efficient Ranking and Selection in Parallel Computing Environments". *Operations Research* 65(3):821–836.

Paulson, E. 1964. "A Sequential Procedure for Selecting the Population with the Largest Mean from k Normal Populations". *The Annals of Mathematical Statistics* 35(1):174–180.

Perng, S. 1969. "A Comparison of the Asymptotic Expected Sample Sizes of Two Sequential Procedures for Ranking Problem". *The Annals of Mathematical Statistics* 40(6):2198–2202.

Rinott, Y. 1978. "On Two-stage Selection Procedures and Related Probability-inequalities". *Communications in Statistics-Theory and Methods* 7(8):799–811.

AUTHOR BIOGRAPHIES

YING ZHONG is a Ph.D. candidate with Hong Kong Ph.D. Fellowship in Department of Management Sciences in College of Business at City University of Hong Kong. His research interests include ranking-and-selection and multi-armed bandit problems with covariates. His email address is yzhong4-c@my.cityu.edu.hk.

L. JEFF HONG is a Chair Professor of Management Sciences in College of Business at City University of Hong Kong. His research interests include simulation modeling and optimization, stochastic programming, and business analytics. He is currently the simulation area editor of *Operations Research* and an associate editor of *Management Science* and *ACM Transactions on Modeling and Computer Simulation*. His email address is jeffhong@cityu.edu.hk.