

Finding probably best systems quickly via simulations

Takayuki Osogami
IBM Tokyo Research Laboratory

We propose an indifference-zone approach for a ranking and selection problem with the goal of reducing both the number of simulated samples of the performance and the frequency of configuration changes. We prove that with a pre-specified high probability our algorithm finds the best system configuration. Our proof hinges on several ideas, including the use of Anderson's probability bound, that have not been fully investigated for the ranking and selection problem. Numerical experiments show that our algorithm can select the best system configuration using up to 50% fewer simulated samples than existing algorithms without increasing the frequency of configuration changes.

Categories and Subject Descriptors: I.6.1 [Simulation and Modeling]: Simulation Theory

General Terms: Theory

Additional Key Words and Phrases: Simulation output analysis, ranking and selection, switching, two-stage, Brownian motion, Indifference zone

1. INTRODUCTION

Computer systems such as Web servers have many possible configurations, and designing a high-performance system requires selecting the best system configuration so that its performance is optimized (e.g., mean response time is minimized). Although the performance of simple systems can be evaluated analytically, the performance of complex systems can only be estimated via computer simulations, or from measurements of real systems. Unfortunately, longer simulation times are often required to estimate the performance more precisely. As a result, it might be computationally intractable to estimate the performance of all system configurations precisely via simulations so as to find the best system configuration.

Our goal is to design an algorithm for finding the best system configuration, so that the total simulation time for the performance estimation is minimized. It is also desirable that the algorithm does not frequently change the system configurations to be simulated, even though frequent configuration changes can reduce the total simulation time. For example, the performance at steady state is measured only after a warm-up period. If we stop simulating a system configuration, π , to simulate other system configurations, we would need to waste another warm-up period before

Author's address: T. Osogami, IBM Tokyo Research Laboratory, 1623-14 Shimotsuruma, Yamato-shi, Kanagawa 242-8502, Japan; email: osogami@jp.ibm.com.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2009 ACM 1049-3301/2009/06-ART12\$10.00

we resume simulating π unless the system state was checkpointed when we stopped simulating π . Therefore, frequent configuration changes result in long warm-up periods or significantly increase the complexity of programming the algorithm for checkpointing.

The problem of selecting the best system configuration, where the system performance is estimated via simulation, has been studied as a ranking and selection problem in the literature [Bechhofer et al. 1995; Swisher et al. 2003]. In the study of ranking and selection, it is standard to assume that simulating the i -th configuration, π_i , yields an independent and identically distributed (i.i.d.) sample from a normal distribution with mean μ_i and standard deviation σ_i for each π_i (both μ_i and σ_i are unknown). Although this assumption may sound unrealistic, it turns out that samples in most simulations can be made arbitrarily close to independent normal random variables by appropriate batching (see Section 9.5.3 of [Law 2006]). For example, consider the mean response time in a queueing system (such as an M/M/1 queue). Although the simulated response time of the r -th job, X_r , does *not* have a normal distribution, the average of the m -th batch of n response times, $\bar{X}_m = \frac{1}{n} \sum_{r=1}^n X_{mn+r}$, approaches a normal random variable by letting n become large for any m . In addition, \bar{X}_m and $\bar{X}_{m'}$ become independent as $n \rightarrow \infty$ for $m \neq m'$. Thus, the mean response time is estimated as $\frac{1}{N/n} \sum_{m=1}^{N/n} \bar{X}_m$ using N total samples, where the batches, $\{\bar{X}_m : m = 0, 1, 2, \dots\}$, are well approximated by a sequence of independent normal random variables.

A goal of ranking and selection algorithms is to find, with a pre-specified high probability ($\geq 1 - \alpha$ for a given α), the configuration π_b that has the largest mean performance μ_b , so that the total number of samples is minimized. We assume that a larger mean performance is better in the rest of this paper. It is common to accept an error within an *indifference-zone* parameter δ , so that any configuration whose mean performance is $> \mu_b - \delta$ is considered to be one of the “best” configurations.

Dudewicz and Dalal [1975] and Rinott [1978] propose two-stage algorithms for finding the best configurations. (The work earlier than [Dudewicz and Dalal 1975; Rinott 1978] assumes known or common σ_i , which is unrealistic for the performance of computer systems.) Stage 1 collects a fixed number of samples from each configuration and estimates the variance. Based on the estimated variance, the total number of samples needed is determined for each configuration, and precisely this number of samples is collected from each configuration in Stage 2. At the end of Stage 2, the configuration that is estimated to be the best is selected (e.g., the one having the largest sample mean). Nelson et al. [2001] propose a screening procedure that can be used with the two-stage algorithms in [Dudewicz and Dalal 1975; Rinott 1978]. The screening procedure eliminates clearly poor configurations after Stage 1, and this can significantly reduce the number of samples needed.

The two-stage algorithm with screening can be extended to more than two stages or to a sequential-stage algorithm, where a screening procedure is applied at each stage until only the best configuration is left [Kim and Nelson 2001]. Although sequential-stage algorithms can reduce the total number of samples, they can require significant overhead and/or increase the complexity of implementing the algorithm, since sequential-stage algorithms can change the configurations to be simulated for *every* sample.

To balance the number of samples and the number of configuration changes, Hong and Nelson [2005] propose a two-stage algorithm that makes decisions sequentially, which we refer to as the HN algorithm. HN interrupts simulation at most once for each system configuration; that is, the algorithm consists of two stages (two-stage algorithm). Also, HN stops sampling from a configuration at any time during Stage 2 when sufficient samples have been collected to make correct decisions; that is, the algorithm makes decisions sequentially (sequential-decision algorithm). Due to this sequential decision-making, HN uses far fewer samples than classical two-stage algorithms such as [Dudewicz and Dalal 1975; Rinott 1978].

The algorithms in [Dudewicz and Dalal 1975; Hong and Nelson 2005; Kim and Nelson 2001; Nelson et al. 2001; Rinott 1978] are called indifference-zone approaches, since their formulation is characterized by the indifference-zone parameter, δ . There are other formulations for selecting the best system configuration via simulation. For example, Chen et al. [2000] and Chen et al. [2000] study a Bayesian approach, where the goal is to maximize the probability that the best configuration is selected under the constraint that the total number of samples is fixed. Chick and Inoue [2001a; 2001b] study other Bayesian approaches, where the goal is to minimize the expected loss, where loss is defined to be the difference between the performances of the selected configuration and the best configuration. Unfortunately, these Bayesian approaches do not guarantee with a pre-specified high probability that the performance of the selected configuration is within δ of the best configuration.

The primary contribution of this paper is a two-stage sequential-decision (TSSD) algorithm that is based on HN with small but essential modifications. In particular, we use a probability bound proved by Anderson (Corollary 4.4 from [Anderson 1960]) for determining a sufficient number of samples to make a correct selection, where HN uses Fabian's probability bound (Theorem 2.3 from [Fabian 1974]). Other small modifications to HN and proof techniques are introduced to prove that TSSD selects the best configuration with a pre-specified high probability.

A secondary contribution of this paper is the numerical evaluation of TSSD. We find that TSSD uses up to 50% fewer samples than HN and up to 10% more samples than KN, depending on the parameter settings. TSSD also compares favorably with a modified KN, where the configurations are changed for every $m \geq 2$ samples.

The rest of the paper is organized as follows. We present the TSSD algorithm in Section 2. In Section 3, we show the effectiveness of TSSD via numerical experiments. Proofs are found in the electronic appendix. In Appendix A, we prove the properties of TSSD for the case where a single configuration has the largest mean performance and the mean performances of the other configurations are at least δ separated from the largest mean performance. In Appendix B, we discuss the general case.

2. THE TSSD ALGORITHM AND ITS PROPERTIES

TSSD selects, with a pre-specified high probability, one of the best configurations from the set of candidate configurations, $\mathcal{C} = \{\pi_1, \dots, \pi_k\}$. A sample from configuration π_i is assumed to have a normal distribution with mean μ_i and standard deviation σ_i for each $\pi_i \in \mathcal{C}$. Let π_b be one of the configurations having the largest

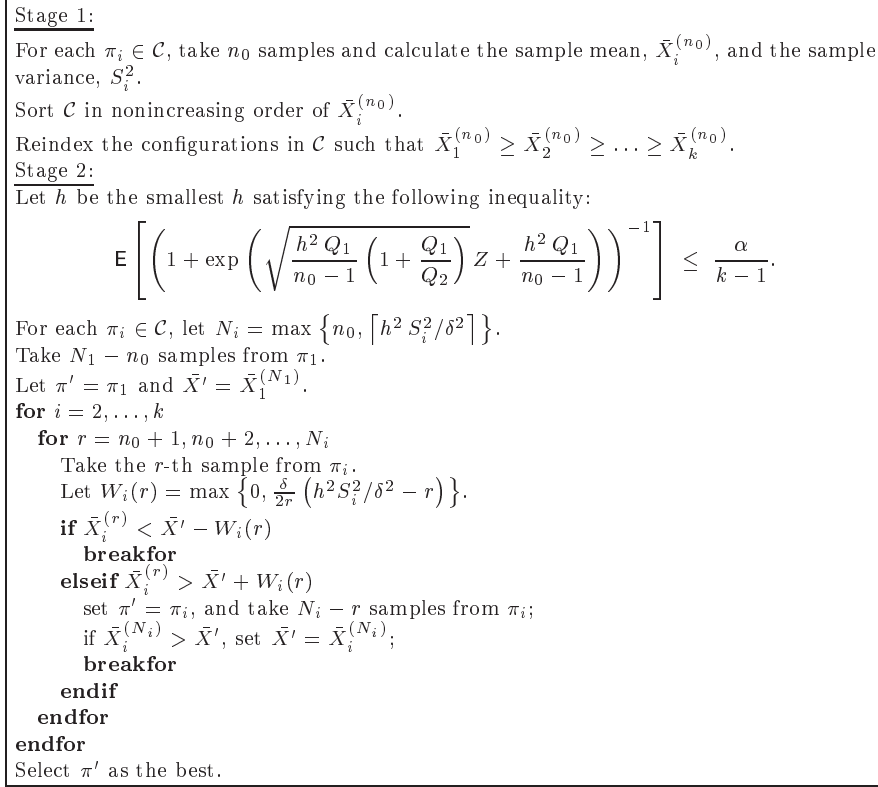


Fig. 1. The TSSD algorithm.

mean such that $\mu_b \geq \mu_i, \forall \pi_i \in \mathcal{C}$. Let $\mathcal{B} = \{\pi_i \in \mathcal{C} \mid \mu_i > \mu_b - \delta\}$ be the set of the best or nearly best configurations, and let $\mathcal{W} = \mathcal{C} \setminus \mathcal{B}$. Formally, the following theorem characterizes the properties of TSSD:

THEOREM 2.1. *Let π be the configuration selected by TSSD. Then $\Pr(\pi \in \mathcal{B}) \geq 1 - \alpha$.*

In this section, we introduce TSSD. The proof of the theorem is postponed to Appendix A and Appendix B. Throughout, we assume that a sample from π_i and a sample from π_j are independent for $\pi_i \neq \pi_j$. The use of “common random numbers” [Law 2006] in TSSD is left as future work.

2.1 The TSSD Algorithm

TSSD consists of two stages. The primary goal of Stage 1 is to estimate the variance of each $\pi_i \in \mathcal{C}$, which will be used in Stage 2 to determine an *upper bound*, N_i , on the number of samples to be collected from each $\pi_i \in \mathcal{C}$. The final goal of Stage 2 is to select the best configuration. Figure 1 summarizes TSSD, which we will elaborate on below. For each $\pi_i \in \mathcal{C}$, $X_{i,r}$ denotes the r -th sample from π_i , and $\bar{X}_i^{(r)} = \frac{1}{r} \sum_{\ell=1}^r X_{i,\ell}$ denotes the sample mean of the first r samples from π_i .

Stage 1 collects a fixed number, n_0 , of samples from each $\pi_i \in \mathcal{C}$. Then the

sample mean, $\bar{X}_i^{(n_0)}$, and the sample variance,

$$S_i^2 = \frac{1}{n_0 - 1} \sum_{\ell=1}^{n_0} \left(X_{i,\ell} - \bar{X}_i^{(n_0)} \right)^2,$$

are calculated. In theory, n_0 can be any integer greater than 1. However, many samples are needed in Stage 1 when n_0 is set large. Also, many samples are needed in Stage 2 when n_0 is set too small, as will be suggested by the arguments below. We set $n_0 = 10$ for our experiments in Section 3, but the suitable n_0 will vary depending on the properties of \mathcal{C} .

At the end of Stage 1, we sort the configurations \mathcal{C} in decreasing order of their sample means. Below, the configurations are re-indexed so that $\bar{X}_i^{(n_0)} \geq \bar{X}_{i+1}^{(n_0)}$ for $1 \leq i \leq k-1$. The sorting often reduces the number of samples to be collected in Stage 2, as we will see in Section 3.2. We remark that Stage 1 of TSSD is the same as Stage 1 of HN except that TSSD calculates S_i^2 for each $\pi_i \in \mathcal{C}$, while HN calculates the sample variance of the differences between the samples for each pair of configurations in \mathcal{C} .

Stage 2 first determines an upper bound on the number of samples to be collected from π_i using

$$N_i = \max \{ n_0, \lceil h^2 S_i^2 / \delta^2 \rceil \} \quad (1)$$

for each $\pi_i \in \mathcal{C}$. Here, h is a confidence parameter, which is defined as the smallest h satisfying

$$\mathbb{E} \left[\frac{1}{1 + \exp \left(\sqrt{\frac{h^2 Q_1}{n_0 - 1} \left(1 + \frac{Q_1}{Q_2} \right)} Z + \frac{h^2 Q_1}{n_0 - 1} \right)} \right] \leq \frac{\alpha}{k-1}, \quad (2)$$

where Z is a standard normal random variable, Q_1 and Q_2 are χ^2 random variables with $n_0 - 1$ degrees of freedom, and the three random variables are independent. Further explanation on the confidence parameter h and an algorithm for calculating h is postponed to Section 2.2.

Now we collect $N_1 - n_0$ additional samples from the (re-ordered) configuration π_1 , and calculate the sample mean of the N_1 samples, $\bar{X}_1^{(N_1)}$. Let $\pi' = \pi_1$ and $\bar{X}' = \bar{X}_1^{(N_1)}$, where π' denotes the provisional best, and \bar{X}' denotes the largest sample mean (so far) of the configurations that have become provisionally best. As we will see below, the provisional best may not have the largest sample mean so far, so \bar{X}' may be larger than $\bar{X}_i^{(N_i)}$ when π_i is the provisional best, π' . Roughly speaking, TSSD determines whether π_i is provisionally best before collecting N_i samples from π_i ; if π is provisionally best, N_i samples are collected from π_i , but $\bar{X}_i^{(N_i)}$ might not be the largest sample mean. We choose to define π' and \bar{X}' in the above way so that we can prove Theorem 2.1.

The rest of the configurations are then sampled in the order determined at the end of Stage 1, but they may or may not be sampled to the upper bounds determined via Equation (1). Note that all of the samples from π_i are taken before a sample is taken from π_{i+1} in Stage 2, for $i = 1, \dots, k-1$. After we take the r -th sample (the

r samples include the n_0 samples collected in Stage 1) from π_i , a decision threshold, $W_i(r)$, is calculated as follows:

$$W_i(r) = \max \left\{ 0, \frac{\delta}{2r} \left(\frac{h^2 S_i^2}{\delta^2} - r \right) \right\}. \quad (3)$$

TSSD takes three different actions based on how $\bar{X}_i^{(r)}$ compares to \bar{X}' .

Case (i): If $\bar{X}_i^{(r)}$ is significantly smaller than \bar{X}' , so that

$$\bar{X}_i^{(r)} < \bar{X}' - W_i(r), \quad (4)$$

then we decide that π_i is not the best configuration and stop sampling from π_i (i.e., the configuration whose sample mean is \bar{X}' eliminates π_i). In addition, if $\pi_i = \pi_k$, we select the provisional best, π' , as the best configuration; otherwise, we start sampling from the next configuration, π_{i+1} .

Case (ii): If $\bar{X}_i^{(r)}$ is significantly larger than \bar{X}' , so that

$$\bar{X}_i^{(r)} > \bar{X}' + W_i(r), \quad (5)$$

then we decide that π_i is better than π' . We set $\pi' = \pi_i$ and take samples from π_i to the upper bound, N_i . If $\bar{X}_i^{(N_i)} \geq \bar{X}'$, we update $\bar{X}' = \bar{X}_i^{(N_i)}$. Note that we do not update \bar{X}' if $\bar{X}_i^{(N_i)} < \bar{X}'$, since \bar{X}' is defined to be the largest sample mean of the configurations that have become provisionally best. In addition, if $\pi_i = \pi_k$, we select π' as the best configuration; otherwise, we start sampling from the next configuration, π_{i+1} .

Case (iii): If $\bar{X}_i^{(r)}$ and \bar{X}' do not have a significant difference, so that

$$\bar{X}' - W_i(r) \leq \bar{X}_i^{(r)} \leq \bar{X}' + W_i(r),$$

then r samples are not sufficient to differentiate π_i and π' , and we continue sampling from π_i . Note that either Inequality (4) or Inequality (5) holds when $r = N_i$, since $W_i(N_i) = 0$.

Stage 2 of TSSD is similar to Stage 2 of HN but differs in several ways. TSSD calculates an upper bound, N_i , on the number of samples for each $\pi_i \in \mathcal{C}$, while HN calculates an upper bound for each pair of configurations. The fundamental difference between TSSD and HN that allows TSSD to use fewer samples than HN is that the upper bounds used in TSSD are determined with the confidence parameter h , while the upper bounds used in HN are determined with a different constant based on Fabian's theorem (Theorem 2.3 from [Fabian 1974]). Also, TSSD uses \bar{X}' in deciding whether the configuration under evaluation becomes provisionally best (Inequalities (4)-(5)), while HN uses a weighted average of the sample mean in Stage 1 and the sample mean in Stage 2 of the provisional best (Equation (3) in [Hong and Nelson 2005]). Finally, \bar{X}' never decreases in TSSD, while the corresponding quantity in HN may decrease. Although the last difference is insignificant with respect to the behavior of the algorithms, we find the difference essential in our proof of Theorem 2.1.

2.2 Decision threshold and confidence parameter

Figure 2 provides further intuition behind the choice of the decision threshold, $W_i(r)$, and the choice of the confidence parameter, h . Observe that $[-r W_i(r), r W_i(r)]$




Fig. 2. The continuation region: TSSD decides whether π_i is better or worse than π_j when $r(\bar{X}_i^{(r)} - \bar{X}')$ exits the continuation region, where $\bar{X}' = \bar{X}_j^{(N_j)}$.

for $0 \leq r \leq N_i$ defines a triangular area, which we refer to as the continuation region. TSSD decides whether π_i is better or worse than the provisional best when $r(\bar{X}_i^{(r)} - \bar{X}')$ exits the continuation region. Specifically, if $r(\bar{X}_i^{(r)} - \bar{X}')$ exits the continuation region through the lower (respectively, upper) edge, then TSSD decides that π_i is worse (respectively, better) than the provisional best.

We need to choose a continuation region such that TSSD makes a correct selection with a pre-specified high probability, and yet the number of samples needed to make decisions (i.e., to exit the continuation region) is minimized. The continuation region depends on two parameters, δ and h , and on a random variable, S_i^2 . Since S_i^2 cannot be controlled, only the confidence parameter h can be used to determine the size of the continuation region for a given indifference-zone parameter, δ . Observe that h does not change the *shape* of the continuation region. A larger h implies a larger continuation region, which in turn forces TSSD to collect stochastically more samples before making a decision but allows TSSD to make a correct decision with higher probability (see Lemma 4 in [Hong 2006]).

We choose the confidence parameter as the smallest h with which we can prove Theorem 2.1 using Inequality (2). Our proof relies on Anderson's results [Anderson 1960] (see also Figure 3):

LEMMA 2.2 COROLLARY 4.4 FROM [ANDERSON 1960]. *Let $a > 0$ and $c > 0$. Let $L(r) = -c + ar$ and $U(r) = c - ar$ for $0 \leq r \leq c/a$, and let $L(r) = U(r) = 0$ for $r > c/a$. Let $B_d(r)$ be a Brownian motion with drift d , where $-\infty < d < \infty$. Let \mathcal{E} be the event that $B_d(r)$ crosses the lower edge of a triangular area, $L(r)$, before crossing the upper edge, $U(r)$. Formally, \mathcal{E} is the event that there exists an r such that $B_d(r) < L(r)$ and $r < s$ for any s such that $B_d(s) > U(s)$. Then,*

$$\Pr(\mathcal{E}) = \mathbf{E} \left[\frac{1}{1 + \exp(2(\sqrt{ac}Z + cd))} \right],$$

where Z is a standard normal random variable.

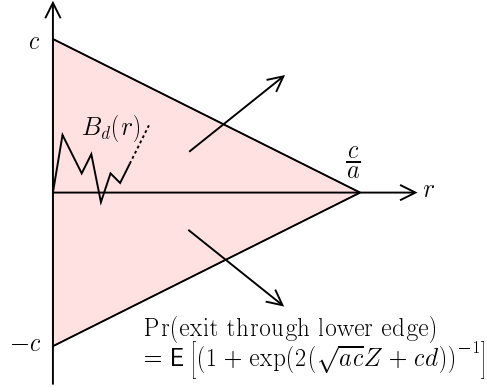


Fig. 3. The probability that a Brownian motion with drift d exits a triangular area through the lower edge, where Z is a standard normal random variable.

Notice the similarities and differences between Figure 2 and Figure 3. For example, the two figures have stochastic processes that start at the origin and eventually exit the triangular areas. However, while $B_d(r)$ in Figure 3 is a Brownian motion with known drift d , the quantity $r(\bar{X}_i^{(r)} - \bar{X}')$ in Figure 2 is a “random walk” with *stochastic* drift $(\mu_i - \bar{X}')$ and has *unknown* parameters, μ_i , σ_i , μ' , and σ' , where μ' and σ' are the mean and standard deviation, respectively, of \bar{X}' . Our proof of Theorem 2.1 will clarify the relationship between Figure 2 and Figure 3.

We calculate the confidence parameter as follows. Let $\eta(h)$ be the left-hand side of Inequality (2). In this paper, we find the smallest h satisfying Inequality (2) by using a binary search, where $\eta(h)$ is evaluated via Monte Carlo integration. Alternatively, we could evaluate $\eta(h)$ via numerical integration,

$$\int_{-\infty}^{\infty} \int_0^{\infty} \int_0^{\infty} \left(1 + \exp \left(\sqrt{\frac{h^2 q_1}{n_0 - 1} \left(1 + \frac{q_1}{q_2} \right)} z + \frac{h^2 q_1}{n_0 - 1} \right) \right)^{-1} d\Psi(q_1) d\Psi(q_2) d\Phi(z),$$

where $\Psi(q)$ is the cumulative distribution function of a χ^2 random variable with $n_0 - 1$ degrees of freedom, and $\Phi(z)$ is the cumulative distribution function of a standard normal random variable. Numerical integration might be more efficient than Monte Carlo integration.

If $\eta(h)$ is a monotonic function of h , a binary search is guaranteed to find the smallest h satisfying Inequality (2). Unfortunately, a proof of the monotonicity of $\eta(h)$ does not appear to be straightforward, even though our numerical experiments suggest $\eta(h)$ is a decreasing function of h . However, it is easy to show that upper and lower bounds of $\eta(h)$ are decreasing in h :

PROPOSITION 2.3. *Let Z be a standard normal random variable, and let Q_1 and Q_2 be χ^2 random variables with $n_0 - 1$ degrees of freedom, where the three random*

variables are independent. Let $\eta(h)$ be the left-hand side of Inequality (2), and let

$$g(h) = \Pr \left(Z \leq - \frac{\frac{Q_1}{n_0-1}}{\sqrt{\frac{\kappa^2}{4h^4} + \frac{Q_1}{(n_0-1)h^2} \left(1 + \frac{Q_1}{Q_2}\right)}} \right),$$

where $\kappa = 1.702$. Then

$$g(h) - 0.01 < \eta(h) < g(h) + 0.01 \quad (6)$$

for $h > 0$. Also, $g(h)$ is a decreasing function of h .

We postpone the proof of the proposition to Appendix C. Note that any h satisfying Inequality (2) will guarantee Theorem 2.1, although the smallest h minimizes the number of samples needed in TSSD. Our binary search is guaranteed to find an h satisfying Inequality (2), even if $\eta(h)$ is not decreasing in h .

The confidence parameter, h , may be compared against the corresponding parameters used in KN and HN. KN determines the upper bound on the number of samples collected from the i -th configuration as

$$N_i^{\text{KN}} = \max \left\{ n_0, \max_{j \neq i} \left\lceil \frac{(h^{\text{KN}})^2 S_{i,j}^2}{2\delta^2} \right\rceil \right\}, \quad (7)$$

where $S_{i,j}^2$ is the sample variance of the difference between samples from π_i and the samples from π_j in Stage 1, and h^{KN} is the confidence parameter for KN. Notice that KN determines N_i^{KN} using $S_{i,j}^2$, while TSSD determines N_i using S_i^2 (see Equation (1)). We define h^{KN} so that the denominator of Equation (7) is $2\delta^2$ instead of δ^2 as in Equation (1), since $\mathbf{E}[S_{i,j}^2] = \mathbf{E}[S_i^2 + S_j^2] = 2\mathbf{E}[S_i^2]$ when the configurations are independent and have a common variance. HN determines the upper bound on the number of samples collected from the i -th configuration as

$$N_i^{\text{HN}} = \max \left\{ n_0, \max_{j \neq i} \left\lceil \frac{(h^{\text{HN}})^2 S_{i,j}^2}{2\delta^2} \right\rceil \right\}, \quad (8)$$

where h^{HN} is the confidence parameter for HN. As with Equation (7), the denominator of Equation (8) is $2\delta^2$ instead of δ^2 .

Figure 4 shows the values of h , h^{KN} , and h^{HN} as functions of α , where $k = 10$ and $n_0 = 10$ are fixed. Specifically, h^{KN} and h^{HN} are calculated using the following equations:

$$h^{\text{KN}} = \sqrt{2(n_0 - 1) \left(\left(\frac{2\alpha}{k-1} \right)^{-\frac{2}{n_0-1}} - 1 \right)} \quad (9)$$

$$h^{\text{HN}} = \sqrt{2(n_0 - 1) \left(\left(2 - 2(1 - \alpha)^{\frac{1}{k-1}} \right)^{-\frac{2}{n_0-1}} - 1 \right)}, \quad (10)$$

so that h^{KN} is the confidence parameter for the fully sequential, indifference-zone procedure with parameter $c = 1$ in [Kim and Nelson 2001], and h^{HN} is the confidence parameter for the minimum switching sequential procedure with Fabian's




Fig. 4. The confidence parameters for HN, KN, and TSSD for given α , where $k = 10$ and $n_0 = 10$.

probability bound using parameter $\lambda = \delta/2$ in [Hong and Nelson 2005]. Observe that h^{KN} and h^{HN} have similar values for all α in Figure 4. This is also suggested by Equations (9)-(10), since $2 - 2(1 - \alpha)^{1/(k-1)} \approx 2\alpha/(k-1)$ when $\alpha \ll 1$. We find that h is 14-18% smaller than h^{KN} and h^{HN} .

3. EXPERIMENTAL RESULTS

In this section, we present and discuss the effectiveness of TSSD with respect to the total number of samples needed, the realized probability of correct selection, and the frequency of changes in the system configurations to be simulated. Specifically, we compare TSSD to the two-stage sequential-decision algorithm by Hong and Nelson [2005] (HN) and the sequential-stage algorithm by Kim and Nelson [2001] (KN). We also compare TSSD against a modification of KN. Recall that KN was not designed with a goal of minimizing the frequency of configuration changes and that KN changes the configurations to be simulated after every sample. The modified KN changes the configurations after every m samples to reduce the frequency of configuration changes. Hong and Nelson [2005] introduce two versions of the algorithm, and the version of HN discussed in this paper is the Minimum Switching Sequential (MSS) procedure with Fabian's probability bound, which is superior to the other version of MSS with Paulson's probability bound.

Throughout, we assume that each configuration generates samples that are i.i.d. according to a normal distribution, where each configuration may have a different mean and/or variance. Also, we set the indifference-zone parameter to $\delta = 1$, the number of samples collected in the first stage $n_0 = 10$, and the number of candidates $k = 10$. The upper bound on the overall error probability is set to $\alpha = 0.05$, but we find that the results of the experiments with $\alpha = 0.1$ are qualitatively similar to those with $\alpha = 0.05$. In this setting, the confidence parameter of TSSD is $h = 4.69$ when $\alpha = 0.05$, and $h = 4.16$ when $\alpha = 0.1$. In all subsequent figures, we take the average of the results from 10,000 runs to generate each data point, since the number of samples, the realized probability of correct selection, and the number of

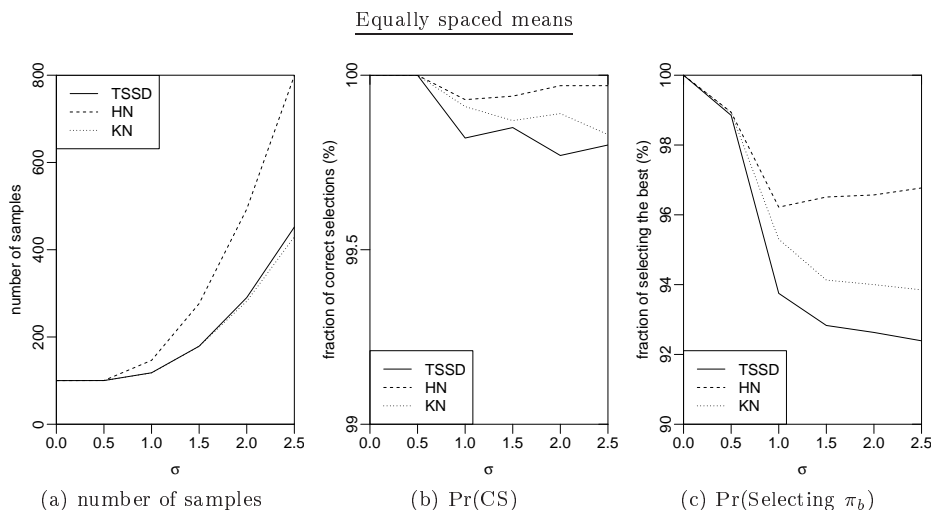


Fig. 5. Comparison of (a) the number of samples, (b) the probability of correct selection, and (c) the probability of selecting π_b in HN (dashed line), KN (dotted line), and TSSD (solid line) when $\alpha = 0.05$. The 10 configurations have a common variance, and their means are equally spaced.

configuration changes are random.

3.1 The number of samples and the probability of correct selection

We start by evaluating the total number of samples needed and the probability of correct selection in HN, KN, and TSSD. In Figure 5, the means of the 10 candidates are chosen between 0.5 and 5 so that they are equally spaced, and the 10 configurations have the same standard deviation, σ , which is varied as indicated by the horizontal axis. Specifically, the i -th smallest mean is $i/2$ for $1 \leq i \leq 10$.

Figure 5(a) plots the number of samples used to select the best configuration for each of HN, KN, and TSSD. When the variance of the performance is small ($\sigma \leq 0.5$), the samples collected in Stage 1 suffice to determine the best configuration, and the number of samples used is $n_0 \times k = 100$ for all three algorithms. As the variability increases, the number of samples increases in all three algorithms. HN has the highest rate of increase, and requires 75% more samples than TSSD for $\sigma = 2.5$. KN has the lowest rate of increase, and requires 8% fewer samples than TSSD for $\sigma = 2.5$.

Figure 5(b) plots the fraction of runs that made correct selections for HN, KN, and TSSD. The settings of the experiments are the same as in Figure 5(a). In particular, recall that $\delta = 1$, so that the two configurations with $\mu_i \geq 4.5$ are in \mathcal{B} . When either of the configurations in \mathcal{B} is selected, it is regarded as a correct selection. Observe that the fraction of correct selections is well above the target of 95% ($\alpha = 0.05$) for all of the three algorithms. Taking a closer look, we find that HN makes correct selections more frequently than TSSD. This makes intuitive sense, since TSSD and HN follow similar procedures but TSSD uses fewer samples than HN before selecting the best configuration. The upper bounds of the number of samples in TSSD are determined based on an analysis that results in an actual

Slippage-configuration means




Fig. 6. Comparison of (a) the number of samples and (b) the probability of correct selection in HN (dashed line), KN (dotted line), and TSSD (solid line) when $\alpha = 0.05$. The 10 configurations have a common variance, and their means are chosen such that $\mu_i = \mu_b - \delta, \forall \pi_i \neq \pi_b$.

probability of the correct selection that is closer to the target than for HN. We also find that HN makes correct selections more frequently than KN. The upper bounds of the numbers of samples in KN and HN are determined using the same (Fabian's) inequality. Since KN makes decisions sequentially, KN is more likely to make a selection with fewer samples and makes incorrect selections more frequently than HN. We find that KN makes correct selections more frequently than TSSD in Figure 5(b), but this is not always the case, as we will see below.

Figure 5(c) plots the fraction of runs that selected π_b for HN, KN, and TSSD. The settings of the experiments are the same as in Figure 5(a) and Figure 5(b). The probability that HN selects π_b is higher than for KN or TSSD. The probability that π_b is selected when $\sigma \leq 0.5$ is significantly higher than that when $\sigma \geq 1.0$ for each of the three algorithms. This is because the samples collected in Stage 1 are more than enough to make correct selections with the pre-specified high probability when $\sigma \leq 0.5$.

In Figure 6, the means of the 10 candidates are chosen such that $\mu_b = \mu_i + \delta, \forall \pi_i \neq \pi_b$. Again, the 10 configurations have the same standard deviation, σ , which varies as indicated by the horizontal axis. This setting of μ_i 's is referred to as a slippage configuration in the literature. Note that $\mathcal{B} = \{\pi_b\}$, and that the mean of every configuration in \mathcal{W} is exactly δ away from μ_b .

Figure 6(a) shows that the three algorithms require more samples for the slippage-configuration means than for the equally spaced means. However, the advantage of TSSD against HN and the disadvantage of TSSD against KN, with respect to the total number of samples, do not change significantly from those for the equally

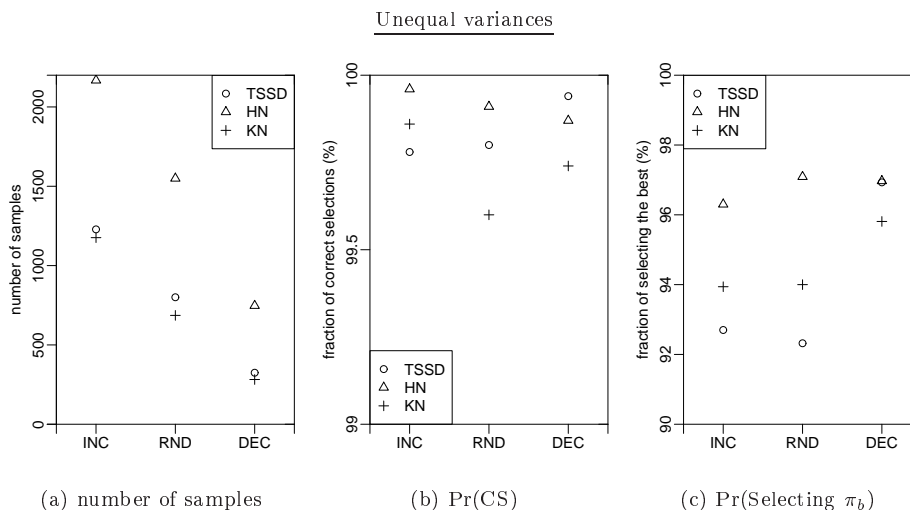


Fig. 7. Comparison of (a) the number of samples, (b) the probability of correct selection, and (c) the probability of selecting π_b in HN (Δ), KN (+), and TSSD (\circ) when $\alpha = 0.05$. The variances of the 10 configurations are not equal, but the means of the 10 configurations are equally spaced.

spaced means. HN uses 69% more samples than TSSD for $\sigma = 2.5$, and KN uses 8% fewer samples than TSSD.

Figure 6(b) plots the the fraction of runs that made correct selections for HN, KN, and TSSD. The settings of the experiments are the same as in Figure 6(a). Since $\mathcal{B} = \{\pi_b\}$ in this setting, the fraction of correct selections is equivalent to the fraction of selecting π_b , so that the figure corresponding to Figure 5(c) is not shown. The fraction of correct selections is closer to the target of 95% for each of the three algorithms than that for the equally spaced means (Figure 5(b)). Again, HN makes correct selections more frequently than KN and TSSD. Unlike Figure 5(b), TSSD makes correct selections more frequently than KN for high σ_i 's.

Figure 7 considers the settings when the 10 configurations have different variances. The means of the 10 configurations are equally spaced, as in Figure 5. We consider three settings: INC, DEC, and RND. Under INC, the standard deviation of a configuration is the same as the mean of the configuration, i.e., $\sigma_i = \mu_i$ for $1 \leq i \leq 10$. Under DEC, the standard deviations are reassigned so that a configuration having a smaller mean has a larger variance, i.e., $\sigma_i = 5.5 - \mu_i$ for $1 \leq i \leq 10$. Under RND, the standard deviation of a configuration is chosen randomly according to $\sigma_i = \mu_{[i]}$ for $1 \leq i \leq 10$, where $[i]$ denotes the i -th element of a random permutation of $\{1, 2, \dots, 10\}$.

Figure 7(a) shows the total number of samples used to select the best configuration for each of the three algorithms under the three settings of the standard deviations, as indicated by the horizontal axis. All of the three algorithms tend to require more samples when the better configurations (which have larger means) have larger variances. This makes intuitive sense since configurations in \mathcal{B} are more likely to be simulated with high precision, and more samples are needed if the configurations in \mathcal{B} have high variances. Taking a close look, we see that HN uses

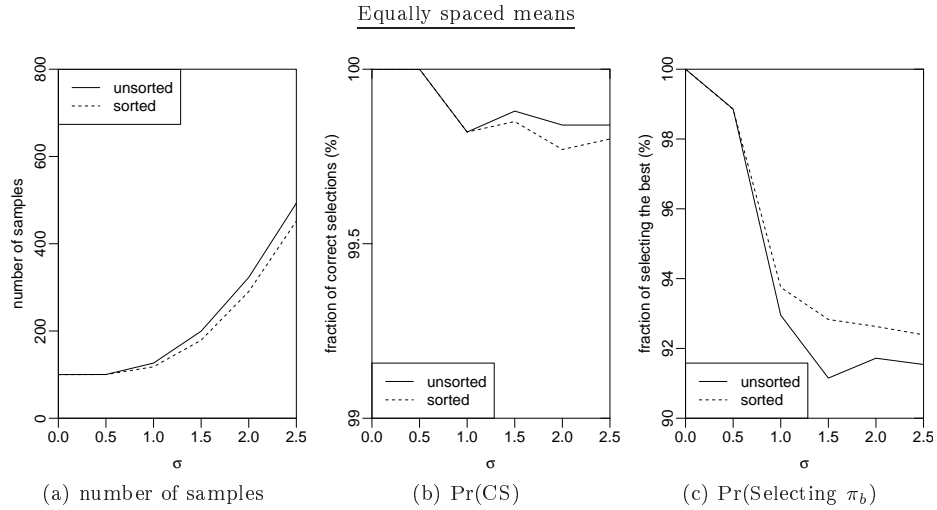


Fig. 8. Comparison of (a) the number of samples, (b) the probability of correct selection, and (c) the probability of selecting π_b in TSSD (dashed line) and a modified TSSD that does not sort configurations before Stage 2 (solid line). The settings are the same as in Figure 5.

75-130% more samples and KN uses 8-14% fewer samples than TSSD.

Figure 7(b) shows the fraction of runs that made correct selections, and Figure 7(c) shows the fraction of runs that selected π_b for HN, KN, and TSSD. The settings of the experiments are the same as in Figure 7(a), so that the two configurations with $\mu_i \geq 4.5$ are in \mathcal{B} . Similarly to the case with common variance (Figure 5), the fraction of correct selections is far above the target of 95% for all of the three algorithms. Also, HN makes correct selections more frequently than KN and TSSD when the variances are chosen as INC or RND.

An interesting observation is that TSSD makes correct selections more frequently than HN when the good configurations have small variances (DEC). This observation may be explained as follows. When the i -th configuration, π_i , is evaluated, TSSD determines the size of the continuation region using S_i^2 , while HN determines the size of the continuation region using S_{ij}^2 , where S_{ij}^2 is the sample variance, in Stage 1, of the differences between the samples from π_i and the samples from the provisional best, π_j . Therefore, when the configurations have different variances, TSSD tends to collect more samples from configurations having larger variances and fewer samples from those having smaller variances than HN. When the variances are chosen as DEC, poor configurations have larger variances, so that TSSD collects fewer samples from the poor configurations than HN. Since the means of the poor configurations are separated by more than δ from the largest mean in the setting under consideration, TSSD can make correct selections with a higher probability using fewer samples than HN.

3.2 Effects of sorting

Recall that TSSD sorts the configurations at the end of Stage 1 in decreasing order of their sample means. In this section, we study the effect of the sorting on the

Slippage-configuration means

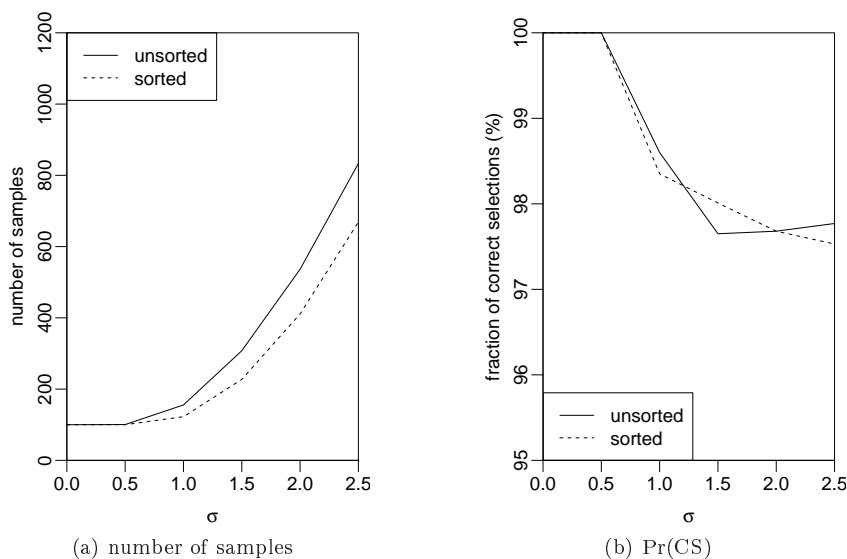


Fig. 9. Comparison of (a) the number of samples and (b) the probability of correct selection in TSSD (dashed line) and in a modified TSSD that does not sort configurations before Stage 2 (solid line). The settings are the same as in Figure 6.

number of samples and the probability of correct selection. Specifically, we evaluate a modified TSSD that operates in the same way as TSSD except that the modified TSSD does not sort the configurations at the end of Stage 1.

Figure 8 evaluates TSSD and the modified TSSD for the same settings as Figure 5, where the means are equally spaced. Figure 8(a) plots the number of samples used to select the best configuration, Figure 8(b) plots the fraction of runs that made correct selections, and Figure 8(c) plots the fraction of runs that selected π_b . Figure 8(a) shows that TSSD reduces the number of samples by up to 11% with sorting. The amount of reduction in the number of samples depends on the standard deviation σ . In Figure 8, the largest reduction is observed when $\sigma = 1.5$. Figure 8(b) and Figure 8(c) show that TSSD reduces the probability of correct selection and the probability of selecting π_b for the settings under consideration, but this is not always the case, as we will see below.

Figure 9 evaluates TSSD and the modified TSSD for the same settings as used for Figure 6, where the μ_i 's have the slippage configuration. Figure 9(a) plots the number of samples collected, and Figure 9(b) plots the fraction of runs that made correct selection, which is equivalent to the fraction of runs that selected π_b for slippage-configuration means. Figure 9(a) shows again that TSSD reduces the number of samples with sorting, but the amount of reduction is larger than for equally spaced means (Figure 8(a)). Specifically, TSSD reduces the number of samples by up to 27%, and the largest reduction is observed when $\sigma = 1.5$. Figure 9(b) shows that sorting may reduce or increase the probability of making correct selections. Overall, we find that sorting has more advantages when the

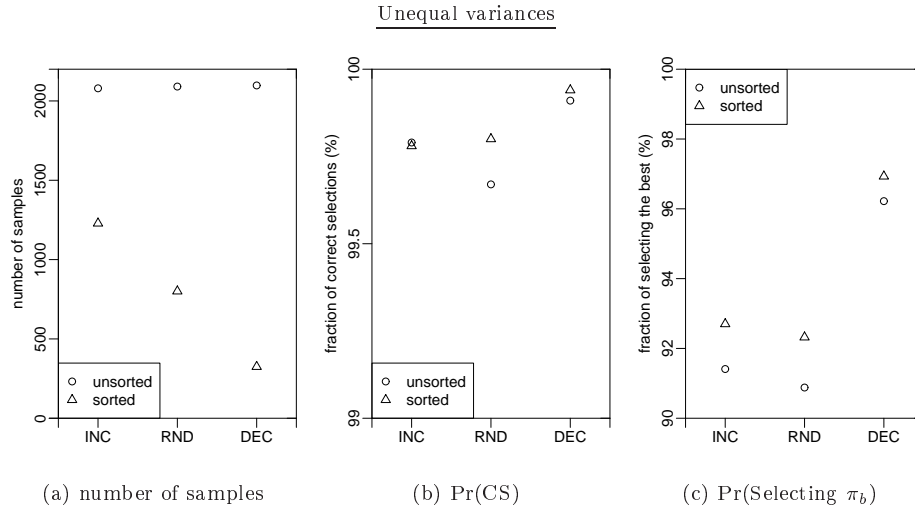


Fig. 10. Comparison of (a) the number of samples, (b) the probability of correct selection, and (c) the probability of selecting π_b in TSSD (Δ) and in a modified TSSD that does not sort configurations before Stage 2 (\circ). The settings are the same as in Figure 7.

means have the slippage configuration than when the means are equally spaced.

Figure 10(a) plots the number of samples needed for TSSD and the modified TSSD. The settings of the experiments are the same as used for Figure 7, where the configurations have different variances. The number of samples is reduced with sorting for all of INC, RND, and DEC, and the largest reduction of 85% is observed when the variances are chosen using DEC. With sorting, TSSD avoids collecting samples from poor configurations up to the upper bounds. For DEC, the poor configurations have large variances and hence require many samples if their samples are collected up to the upper bounds, which results in the large reductions in the numbers of samples with sorting. The effect of sorting in reducing the number of samples diminishes when the poor configurations have smaller variances. In Figure 10(a), sorting reduces the number of samples by 62% for RND and 40% for INC. Overall, the amount of reduction in the number of samples with sorting shown in Figure 10(a) is larger than those shown in Figure 8(a) and Figure 9(a). When some of the configurations have large variances as in Figure 10(a), whether or not the samples from those configurations having large variances are collected up to the upper bounds has a large impact on the numbers of samples.

Figure 10(b) plots the fraction of runs that made correct selections, and Figure 10(b) plots the fraction of runs that selected π_b for TSSD and for the modified TSSD. The settings of the experiments are the same as Figure 10(a). The probabilities of correct selection and of selecting π_b are reduced with sorting for most cases, but the probability of correct selection for INC does not appear to be reduced by sorting.

3.3 Impact of changing configurations

We have seen that TSSD uses significantly fewer samples than HN and that KN uses even fewer samples than TSSD. Since KN was not designed to reduce the frequency of configuration changes, KN changes configurations far more frequently than TSSD. An obvious modification to KN to reduce the frequency of configuration changes is to change configurations after every m samples for $m \geq 2$. However, our experiments (data not shown) suggest that TSSD is superior to the modified KN for any m with respect to at least one of the two metrics: the number of samples or the number of configuration changes.

The cost of changing configurations is sometimes significant. For example, consider a problem of finding an optimal configuration (design) for a complex system such as a car and a factory out of several candidate configurations. Simulating such a complex system often requires huge working-set memory, and changing configurations results in changing the working-set memory. The actual cost of changing configurations varies widely depending on the hardware, the operating system, and the application running the simulations. However, even under ideal conditions, it would take at least 3.3 seconds to transfer a 1 GB working-set memory from a disk to the main memory via Serial Attached SCSI that can transfer up to 300 MB per second. Sequential algorithms that require thousands of configuration changes might become prohibitive when each configuration change requires a few seconds.

4. CONCLUSION

In this paper, we propose a two-stage sequential-decision (TSSD) algorithm that finds, with a pre-specified high probability, one of the best system configurations, so that the total simulation time for performance estimation is minimized. Our numerical experiments show that TSSD improves upon existing algorithms with respect to the total simulation time and the frequency of configuration changes. An additional contribution of this paper is a formal discussion of the statement, “it is more difficult to select π_b when $\mathcal{B} = \{\pi_b\}$ than to select one of \mathcal{B} when $|\mathcal{B}| \geq 2$ for a fixed \mathcal{W} .” Existing proofs of this statement for algorithms that make decisions sequentially do not formally take into account the following property. Under sequential-decision algorithms, the best configuration may be eliminated by a nearly best configuration with high probability, and nearly best configurations may be eliminated by a poor configuration with high probability. Therefore, it is not quite obvious that a sequential-decision algorithm can select one of any nearly best configurations with a pre-specified high probability.

TSSD has applications in optimizing system performance by selecting the system configuration having the best simulated performance. TSSD is also suitable for selecting the best configuration when the system performance is measured via physical experiments instead of computer simulations. Note that changing system configurations may be difficult in physical experiments, but TSSD requires changing the system configurations only infrequently.

In theory, TSSD can be applied to any number of candidate configurations, but in practice TSSD is suitable for a relatively small number of candidate configurations ($1 < k < 20$). When there are more candidate configurations, TSSD may be combined with local search algorithms to selectively evaluate only promising config-

urations. We remark that TSSD and the ideas in TSSD can be used to improve key components of the local search algorithms proposed in [Osogami and Itoko 2006; Pichitlamken and Nelson 2003], specifically the neighborhood search for finding a configuration that is better than the current configuration.

For readability, this paper considers only a single *shape* of the continuation region, but this can be generalized. For example, the slope of the upper and lower boundary of the continuation region is fixed at $\pm\delta/2$ throughout this paper, but it is straightforward to extend the slope to $\pm\delta/(2c)$ for any c in $0 < c < \infty$.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library by visiting the following URL: <http://www.acm.org/pubs/citations/journals/tomacs/2009-19-3/p1-osogami>.

ACKNOWLEDGMENTS

I would like to thank Shannon Jacobs for English rewriting as well as the associate editor and the anonymous referees for their detailed comments on the previous versions of this paper.

REFERENCES

- ANDERSON, T. W. 1960. A modification of the sequential probability ratio test to reduce the sample size. *Annals of Mathematical Statistics* 31, 1, 165–197.
- BECHHOFFER, R. E., SANTNER, T. J., AND GOLDSMAN, D. M. 1995. *Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons*. John Wiley & Sons, New York, NY.
- BOESEL, J., NELSON, B. L., AND KIM, S. 2003. Using ranking and selection to “clean up” after simulation optimization. *Operations Research* 51, 5, 814–825.
- CHEN, C.-H., LIN, J., YÜCESAN, E., AND CHICK, S. E. 2000. Simulation budget allocation for further enhancing the efficiency of ordinal optimization. *Discrete Event Dynamic Systems: Theory and Applications* 10, 3, 251–270.
- CHEN, H.-C., CHEN, C.-H., AND YÜCESAN, E. 2000. Computing efforts allocation for ordinal optimization and discrete event simulation. *IEEE Transactions on Automatic Control* 45, 5, 960–964.
- CHICK, S. E. AND INOUE, K. 2001a. New procedures to select the best simulated system using common random numbers. *Management Science* 47, 8, 1133–1149.
- CHICK, S. E. AND INOUE, K. 2001b. New two-stage and sequential procedures for selecting the best simulated system. *Operations Research* 49, 1, 732–743.
- DUDEWICZ, E. J. AND DALAL, S. R. 1975. Allocation of observations in ranking and selection with unequal variances. *Sankhya B* 37, 1, 28–78.
- FABIAN, V. 1974. Note on Anderson’s sequential procedures with triangular boundary. *The Annals of Statistics* 2, 1, 170–176.
- GRIMMETT, G. AND STIRZAKER, D. 2001. *Probability and Random Processes*, Third ed. Oxford University Press, New York, NY.
- HONG, L. J. 2006. Fully sequential indifference-zone selection procedures with variance-dependent sampling. *Naval Research Logistics* 53, 5, 464–476.
- HONG, L. J. AND NELSON, B. L. 2005. The tradeoff between sampling and switching: New sequential procedures for indifference-zone selection. *IIE Transactions* 37, 7, 623–634.
- JENNISON, C., JOHNSTON, I. M., AND TURNBULL, B. W. 1982. Asymptotically optimal procedures for sequential adaptive selection of the best of several normal means. In *Statistical Decision Theory and Related Topics III*. Vol. 2. Academic Press, New York, NY, 55–86.
- ACM Transactions on Modeling and Computer Simulation, Vol. 19, No. 3, Article No. 12, June 2009.

- KIM, S. AND NELSON, B. L. 2001. A fully sequential procedure for indifference-zone selection in simulation. *ACM Transactions on Modeling and Computer Simulation* 11, 3, 251–273.
- LAW, A. M. 2006. *Simulation Modeling and Analysis*, Fourth ed. McGraw-Hill, New York, NY.
- NELSON, B. L., SWANN, J., GOLDSMAN, D., AND SONG, W. 2001. Simple procedures for selecting the best simulated system when the number of alternatives is large. *Operations Research* 49, 6, 950–963.
- OSOGAMI, T. AND ITOKO, T. 2006. Finding probably better system configurations quickly. In *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems*. ACM, New York, NY, 264–275.
- PICHITLAMKEN, J. AND NELSON, B. 2003. A combined procedure for optimization via simulation. *ACM Transactions on Modeling and Computer Simulation* 13, 2, 155–179.
- RINOTT, Y. 1978. On two-stage selection procedures and related probability-inequalities. *Communications in Statistics — Theory and Methods* A7, 8, 799–811.
- STEIN, C. 1945. A two-sample test for a linear hypothesis whose power is independent of the variance. *The Annals of Mathematical Statistics* 16, 3, 243–258.
- SWISHER, J. R., JACOBSON, S. H., AND YÜCESAN, E. 2003. Discrete-event simulation optimization using ranking, selection, and multiple comparison procedures: Survey. *ACM Transactions on Modeling and Computer Simulation* 13, 2, 134–154.

Received December 2006; revised June 2007, January 2008; accepted June 2008.

THIS DOCUMENT IS THE ONLINE-ONLY APPENDIX TO:

Finding probably best systems quickly via simulations

Takayuki Osogami
IBM Tokyo Research Laboratory

ACM Transactions on Modeling and Computer Simulation, Vol. 19, No. 3, Article No. 12, June 2009.

A. PROOF OF THEOREM 2.1 FOR SOLO BEST CONFIGURATION

In this section, we prove Theorem 2.1 for the case where $|\mathcal{B}| = 1$. Our proof uses Lemma C.1 in Section C, which differs from the proof of HN [Hong and Nelson 2005] that uses Fabian’s theorem. Fabian’s theorem provides a bound on the probability that a random walk with (deterministic) drift exits a triangular area through the lower edge (recall Figure 3), while Lemma C.1 provides a bound for a “random walk” with *stochastic* drift. Lemma C.1 is proved using Lemma 2.2, which provides a bound on the probability that a Brownian motion with (deterministic) drift exits a triangular area through the lower edge. In Section A.1, we prove that TSSD selects π_b with a pre-specified high probability under the condition that configurations are not sorted at the end of Stage 1. The proof will be modified in Section A.2 to cover the case where the configurations are sorted.

A.1 Without sorting

Throughout this section, we assume that $\mathcal{B} = \{\pi_b\}$ and that the configurations are not sorted at the end of Stage 1. We will prove that, with probability $\geq 1 - \alpha$, π_b becomes a provisional best, and $\pi_i \in \mathcal{W}$ does not become a provisional best after π_b becomes a provisional best. Let \mathcal{C}_{bef} and \mathcal{C}_{aft} , respectively, be the sets of configurations that are sampled before and after π_b in Stage 2. Note that $\pi_b \notin \mathcal{C}_{\text{bef}} \cup \mathcal{C}_{\text{aft}}$. Let ICS_i be the event that TSSD makes the incorrect selection that $\pi_i \in \mathcal{W}$ is better than π_b . That is, $\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}})$ is the probability that π_b does not become a provisional best given that $\bar{X}' = \bar{X}_i^{(N_i)}$ immediately before π_b is sampled in Stage 2, and $\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{aft}})$ is the probability that π_i becomes a provisional best given that $\pi' = \pi_b$ immediately before π_i is sampled in Stage 2. Since π_b is selected as the best configuration if no $\pi_i \in \mathcal{W}$ becomes a provisional best, the probability of an incorrect selection, $\Pr(\text{ICS})$, (i.e., π_b is not selected as the best configuration) is bounded as follows:

$$\Pr(\text{ICS}) \leq \sum_{\pi_i \in \mathcal{C}_{\text{bef}}} \Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}}) + \sum_{\pi_i \in \mathcal{C}_{\text{aft}}} \Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{aft}})$$

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2009 ACM 1049-3301/2009/06-ART12\$10.00

We will prove the following two inequalities:

$$\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}}) \leq \frac{\alpha}{k-1} \quad (11)$$

and

$$\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{aft}}) \leq \frac{\alpha}{k-1}. \quad (12)$$

Note that Inequalities (11)-(12) imply $\Pr(\text{ICS}) \leq \alpha$, since $|\mathcal{C}_{\text{bef}}| + |\mathcal{C}_{\text{aft}}| = |\mathcal{W}| = k-1$.

We will first prove Inequality (11). Given that $\bar{X}' = \bar{X}_i^{(N_i)}$ immediately before sampling from π_b in Stage 2, π_b does not become the provisional best only when $\bar{X}_b^{(r)}$ crosses a lower edge, $\bar{X}' - W_b(r)$, before crossing an upper edge, $\bar{X}' + W_b(s)$ (see Figure 2). Formally, ICS_i occurs only when there exists an r such that $\bar{X}_b^{(r)} < \bar{X}' - W_b(r)$ and $r < s$ for any s such that $\bar{X}_b^{(s)} > \bar{X}' + W_b(s)$. Thus,

$$\begin{aligned} & \Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}}) \\ & \leq \Pr\left(\pi_b \text{ does not become } \pi' \mid \bar{X}' = \bar{X}_i^{(N_i)}\right) \\ & = \Pr\left(\exists r \leq N_b \text{ s.t. } \bar{X}_b^{(r)} < \bar{X}' - W_b(r) \text{ and } r < s \forall s \text{ s.t. } \bar{X}_b^{(s)} > \bar{X}' + W_b(s) \right. \\ & \quad \left. \mid \bar{X}' = \bar{X}_i^{(N_i)}\right), \end{aligned} \quad (13)$$

for any $\pi_i \in \mathcal{C}_{\text{bef}}$.

We will now find an upper bound on the right-hand side of Inequality (13), so that Lemma C.1 can be applied to the upper bound. Observe that $\bar{X}_b^{(t)}$ becomes more likely to cross a lower edge, $\bar{X}' - W_b(t)$, before crossing an upper edge, $\bar{X}' + W_b(t)$, if both the lower edge and the upper edge are shifted up. Formally, since $\mu_b > \mu_i + \delta$, we have $\bar{X}_b^{(t)} < \bar{X}' - W_b(t) + (\mu_b - \mu_i - \delta)$ whenever $\bar{X}_b^{(t)} < \bar{X}' - W_b(t)$, and $\bar{X}_b^{(t)} > \bar{X}' + W_b(t) + (\mu_b - \mu_i - \delta)$ implies $\bar{X}_b^{(t)} > \bar{X}' + W_b(t)$ for any t . Hence, Inequality (13) implies that $\pi_i \in \mathcal{C}_{\text{bef}}$ remains the provisional best after sampling from π_b only if $\bar{X}_b^{(r)}$ crosses the shifted lower-edge, $\bar{X}' - W_b(r) + \mu_b - \mu_i - \delta$, before crossing the shifted upper-edge, $\bar{X}' + W_b(s) + \mu_b - \mu_i - \delta$. Therefore,

$$\begin{aligned} & \Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}}) \\ & \leq \Pr\left(\exists r \leq N_b \text{ s.t. } \bar{X}_b^{(r)} < \bar{X}' - W_b(r) + \mu_b - \mu_i - \delta \right. \\ & \quad \left. \text{and } r < s \forall s \text{ s.t. } \bar{X}_b^{(s)} > \bar{X}' + W_b(s) + \mu_b - \mu_i - \delta\right). \end{aligned} \quad (14)$$

The right-hand side of Inequality (14) is the probability that $r(\bar{X}_b^{(r)} - \bar{X}')$ exits the continuation region through the lower edge when the continuation region is shifted up by $(\mu_b - \mu_i - \delta)$, and thus is an upper bound of Inequality (13). Now, let $Z_{b,\ell} = (X_{b,\ell} - \mu_b)/\sigma_b$ for $\ell = 1, 2, \dots$ and $Z_i = -(\bar{X}' - \mu_i)/\sqrt{\sigma_i^2/N_i}$, where we recall that $\bar{X}' = \bar{X}_i^{(N_i)}$. Observe that N_i is a random variable with support on $\{n_0, n_0 + 1, \dots\}$. Since the sample mean and the sample *variance* of a normal distribution are independent (see Section 4.10 of [Grimmett and Stirzaker 2001]), N_i is independent of $\sum_{\ell=1}^{n_0} X_{i,\ell}$ even though N_i is determined using the sample

variance of $X_{i,1}, \dots, X_{i,n_0}$. Therefore, Lemma C.3 in Section C implies that Z_i is a standard normal random variable, and that Z_i and N_i are independent. Also, the $Z_{b,\ell}$'s are standard normal random variables, and the $Z_{b,\ell}$'s and Z_i are independent. Now, observe that, by Equation (3),

$$\begin{aligned} \bar{X}_b^{(t)} &< \bar{X}' - W_b(t) + \mu_b - \mu_i - \delta \\ &\Leftrightarrow \frac{\sigma_b}{t} \sum_{\ell=1}^t Z_{b,\ell} + \sqrt{\frac{\sigma_i^2}{N_i}} Z_i + \delta < -\max \left\{ 0, \frac{\delta}{2t} \left(\frac{h^2 S_b^2}{\delta^2} - t \right) \right\}, \end{aligned}$$

and that

$$\begin{aligned} \bar{X}_b^{(t)} &> \bar{X}' + W_b(t) + \mu_b - \mu_i - \delta \\ &\Leftrightarrow \frac{\sigma_b}{t} \sum_{\ell=1}^t Z_{b,\ell} + \sqrt{\frac{\sigma_i^2}{N_i}} Z_i + \delta > \max \left\{ 0, \frac{\delta}{2t} \left(\frac{h^2 S_b^2}{\delta^2} - t \right) \right\}. \end{aligned}$$

Thus, by Inequality (14), $\pi_i \in \mathcal{C}_{\text{bef}}$ remains the provisional best after sampling from π_b only if

$$R(t) = \sum_{\ell=1}^t \left(Z_{b,\ell} + \frac{1}{\sigma_b} \left(\delta + \sqrt{\sigma_i^2 / N_i} Z_i \right) \right)$$

crosses a lower edge, $-h^2 S_b^2 / (2\delta\sigma_b) + \delta r / (2\sigma_b)$, before crossing an upper edge, $h^2 S_b^2 / (2\delta\sigma_b) - \delta r / (2\sigma_b)$. Therefore,

$$\begin{aligned} &\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}}) \\ &\leq \Pr \left(\exists r \leq N_b \text{ s.t. } R(r) < \min \left\{ 0, -\frac{h^2 S_b^2}{2\delta\sigma_b} + \frac{\delta}{2\sigma_b} r \right\} \right. \\ &\quad \left. \text{and } r < s \forall s \text{ s.t. } R(s) > \max \left\{ 0, \frac{h^2 S_b^2}{2\delta\sigma_b} - \frac{\delta}{2\sigma_b} s \right\} \right). \end{aligned} \quad (15)$$

We will now apply Lemma C.1 in Section C to Inequality (15), and simplify the formulas to obtain Inequality (11). To apply Lemma C.1, we condition on S_b^2 and N_i in Inequality (15). Given S_b^2 and N_i , $c = h^2 S_b^2 / (2\delta\sigma_b)$ and $a = \delta / (2\sigma_b)$ are constants and $(\delta + \sqrt{\sigma_i^2 / N_i} Z_i) / \sigma_b$ is a normal random variable with mean $\delta / \sigma_b \geq 0$. Lemma C.1 now implies

$$\begin{aligned} &\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}}) \\ &\leq \mathbb{E} \left[\mathbb{E} \left[\frac{1}{1 + \exp \left(2 \left(\sqrt{\frac{h^2 S_b^2}{2\delta\sigma_b}} \cdot \frac{\delta}{2\sigma_b} \cdot Z + \frac{h^2 S_b^2}{2\delta\sigma_b} \cdot \frac{\delta}{\sigma_b} \left(1 + \frac{1}{\delta} \sqrt{\frac{\sigma_i^2}{N_i}} Z_i \right) \right) \right)} \middle| S_b^2, N_i \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\frac{1}{1 + \exp \left(\sqrt{\frac{h^2 S_b^2}{\sigma_b^2}} Z + \frac{h^2 S_b^2}{\sigma_b^2} \left(1 + \frac{1}{\delta} \sqrt{\frac{\sigma_i^2}{N_i}} Z_i \right) \right) \right)} \middle| S_b^2, N_i \right] \right], \end{aligned} \quad (16)$$

where Z is a standard normal random variable. Since Z and Z_i are independent by Lemma C.1, Inequality (16) is simplified to

$$\begin{aligned} & \Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}}) \\ & \leq \mathbf{E} \left[\mathbf{E} \left[\frac{1}{1 + \exp \left(\sqrt{\frac{h^2 S_b^2}{\sigma_b^2} + \left(\frac{h^2 S_b^2}{\sigma_b^2} \right)^2 \frac{\sigma_i^2}{\delta^2 N_i}} Z' + \frac{h^2 S_b^2}{\sigma_b^2} \right)} \middle| S_b^2, N_i \right] \right], \end{aligned}$$

where Z' is a standard normal random variable. Note that Z' is independent of S_b^2 and N_i , since Z and Z_i are independent of S_b^2 and N_i . Since $\mathbf{E}[\mathbf{E}[f(\mathbf{A}, \mathbf{B}) \mid \mathbf{A}]] = \mathbf{E}[f(\mathbf{A}, \mathbf{B})]$ for a function, f , when random vectors, \mathbf{A} and \mathbf{B} , are independent, it follows that

$$\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}}) = \mathbf{E} \left[\frac{1}{1 + \exp \left(\sqrt{\frac{h^2 S_b^2}{\sigma_b^2} + \left(\frac{h^2 S_b^2}{\sigma_b^2} \right)^2 \frac{\sigma_i^2}{\delta^2 N_i}} Z' + \frac{h^2 S_b^2}{\sigma_b^2} \right)} \right] \quad (17)$$

By Lemma C.2 in Section C, the expression inside the expectation of Inequality (17) is a nonincreasing function of N_i . Since $N_i \geq h^2 S_i^2 / \delta^2$ from Equation (1), we obtain

$$\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}}) \leq \mathbf{E} \left[\frac{1}{1 + \exp \left(\sqrt{\frac{h^2 S_b^2}{\sigma_b^2} + h^2 \left(\frac{S_b^2}{\sigma_b^2} \right)^2 \frac{\sigma_i^2}{S_i^2}} Z + \frac{h^2 S_b^2}{\sigma_b^2} \right)} \right].$$

Let $Q_b = (n_0 - 1)S_b^2/\sigma_b^2$ and $Q_i = (n_0 - 1)S_i^2/\sigma_i^2$. Then Q_b and Q_i are χ^2 random variables with $n_0 - 1$ degrees of freedom, and Q_b , Q_i , and Z are all independent. Note that Q_i and Z_i are independent due to a property of the sample mean and sample variance of a normal distribution (see Section 4.10 of [Grimmett and Stirzaker 2001]). Therefore,

$$\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}}) \leq \mathbf{E} \left[\frac{1}{1 + \exp \left(\sqrt{\frac{h^2 Q_b}{n_0 - 1} \left(1 + \frac{Q_b}{Q_i} \right)} Z + \frac{h^2 Q_b}{n_0 - 1} \right)} \right]. \quad (18)$$

Now, Inequality (11) follows from the way we choose h (see Inequality (2)). This completes the proof of Inequality (11).

It only remains to prove Inequality (12). Recall that when π_b becomes the provisional best, \bar{X}' is updated to $\bar{X}_b^{(N_b)}$ if $\bar{X}_b^{(N_b)} > \bar{X}'$ and remains unchanged if $\bar{X}_b^{(N_b)} \leq \bar{X}'$. Since \bar{X}' is only updated to a larger value in TSSD, $\bar{X}' \geq \bar{X}_b^{(N_b)}$ always holds after π_b becomes the provisional best. Therefore, given that π_b becomes the provisional best, $\pi_i \in \mathcal{C}_{\text{aft}}$ becomes the provisional best only when $\bar{X}_i^{(r)}$ crosses an upper edge, $\bar{X}' + W_i(r)$, before crossing a lower edge, $\bar{X}' - W_i(s)$, where

$\bar{X}' \geq \bar{X}_b^{(N_b)}$. Formally, the event ICS_i occurs only when there exists an r such that $\bar{X}_i^{(r)} > \bar{X}' + W_i(r)$ and $r < s$ for any s such that $\bar{X}_i^{(s)} < \bar{X}' - W_i(s)$. Thus,

$$\begin{aligned}
 & \Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{aft}}) \\
 & \leq \Pr\left(\pi_i \text{ becomes } \pi' \mid \bar{X}' \geq \bar{X}_b^{(N_b)}\right) \\
 & = \Pr\left(\exists r \leq N_i \text{ s.t. } \bar{X}_i^{(r)} > \bar{X}' + W_i(r) \right. \\
 & \quad \left. \text{and } r < s \forall s \text{ s.t. } \bar{X}_i^{(s)} < \bar{X}' - W_i(s) \mid \bar{X}' \geq \bar{X}_b^{(N_b)}\right) \\
 & \leq \Pr\left(\exists r \leq N_i \text{ s.t. } \bar{X}_i^{(r)} > \bar{X}_b^{(N_b)} + W_i(r) \right. \\
 & \quad \left. \text{and } r < s \forall s \text{ s.t. } \bar{X}_i^{(s)} < \bar{X}_b^{(N_b)} - W_i(s)\right) \tag{19} \\
 & = \Pr\left(\pi_i \text{ becomes } \pi' \mid \bar{X}' = \bar{X}_b^{(N_b)}\right) \\
 & \leq \Pr\left(\exists r \leq N_i \text{ s.t. } -\bar{X}_i^{(r)} < -\bar{X}_b^{(N_b)} - W_i(r) \right. \\
 & \quad \left. \text{and } r < s \forall s \text{ s.t. } -\bar{X}_i^{(s)} > -\bar{X}_b^{(N_b)} + W_i(s)\right).
 \end{aligned}$$

Since $-\mu_i > -\mu_b + \delta$, the last expression is equivalent to Inequality (13) and can be shown to be $\leq \alpha/(k-1)$. This completes the proof of Theorem 2.1 for the case without sorting and $\mathcal{B} = \{\pi_b\}$.

A.2 With sorting

In this section, we prove that sorting the configurations at the end of Stage 1 does not change (the upper bound of) the probability of an incorrect selection. When the configurations are sorted, $\bar{X}_b^{(n_0)} < \bar{X}_i^{(n_0)}$ implies $\pi_i \in \mathcal{C}_{\text{bef}}$. Hence sorting increases $\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}})$, since $r(\bar{X}_b^{(r)} - \bar{X}_i^{(N_i)})$ at $r = n_0$ is more likely to be in the lower region of Figure 2 when $\bar{X}_b^{(n_0)} < \bar{X}_i^{(n_0)}$. However, $\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{aft}})$ is also larger when $\bar{X}_b^{(n_0)} < \bar{X}_i^{(n_0)}$, since $r(\bar{X}_i^{(r)} - \bar{X}_b^{(N_b)})$ at $r = n_0$ is more likely to be in the upper region of Figure 2. Thus, the increase in $\Pr(\text{ICS}_i)$ is not due to sorting but due to the outcomes of the sample means in Stage 1. Before formalizing the above idea, we will briefly discuss related prior work.

Boesel et al. also prove that sorting based on the sample means in Stage 1 does not increase the probability of an incorrect selection for their sort-and-iterative-screen procedure (Section 6.2 of the online companion for [Boesel et al. 2003]). Their proof relies on the fact that sorting does not stochastically change the number of samples collected in Stage 2 for each configuration. This means that $\Pr(\text{ICS} \mid \mathcal{C}_{\text{bef}}) = \Pr(\text{ICS} \mid \mathcal{C}_{\text{aft}})$ holds given any sample means in Stage 1. Since TSSD collects samples from a configuration up to the upper bound if the configuration is found to be the provisional best, sorting may change the number of samples actually collected in Stage 2. We thus cannot prove Theorem 2.1 using the techniques introduced by Boesel et al. [2003].

Jennison et al. [1982] argue that the difference between the sample means of two configurations cannot be embedded on a Brownian motion (with a constant drift) if the sampling rule depends on the sample means of individual configurations. The

sampling rule depends on the sample means of individual configurations when the configurations are sorted. This is why HN uses, instead of \bar{X}' , a weighted average of the sample mean in Stage 1 and the sample mean in Stage 2. In our proof, we embed the difference between the sample means on a Brownian motion with a drift having a normal distribution, so it is in agreement with the results from Jennison et al.

A result of Stein [1945] (page 245) requires that the rule of allocating the number of samples to be collected in Stage 2 should be specified initially and should not depend on the sample means in Stage 1 for two-stage algorithms that do not make decisions sequentially. For TSSD, the result of Stein corresponds to the requirement that the upper bounds on the number of samples collected from a configuration should not depend on the sample means in Stage 1. Although sorting may change the number of samples actually collected in Stage 2, it does not change the upper bound determined with Equation (1). As we will prove below, sorting does not increase the probability of an incorrect selection for TSSD.

Let E_i be the event that π_b is eliminated by $\pi_i \neq \pi_b$ in TSSD (with sorting) given that the two configurations are compared directly with each other. Recall Inequality (13) and Inequality (19), where we show that

$$\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{bef}}) \leq \Pr(\pi_b \text{ does not become } \pi' \mid \bar{X}' = \bar{X}_i^{(N_i)})$$

and

$$\Pr(\text{ICS}_i \mid \pi_i \in \mathcal{C}_{\text{aft}}) \leq \Pr(\pi_i \text{ becomes } \pi' \mid \bar{X}' = \bar{X}_b^{(N_b)}).$$

This argument carries over to the case when the configurations are sorted, so that $\Pr(\text{ICS}_i) \leq \Pr(E_i)$. Observe that, for $\pi_i \neq \pi_b$,

$$\begin{aligned} \Pr(E_i) &= \Pr(E_i \mid \pi_i \in \mathcal{C}_{\text{aft}} \text{ and } \bar{X}_i^{(n_0)} < \bar{X}_b^{(n_0)}) \Pr(\bar{X}_i^{(n_0)} < \bar{X}_b^{(n_0)}) \\ &\quad + \Pr(E_i \mid \pi_i \in \mathcal{C}_{\text{bef}} \text{ and } \bar{X}_i^{(n_0)} > \bar{X}_b^{(n_0)}) \Pr(\bar{X}_i^{(n_0)} > \bar{X}_b^{(n_0)}), \end{aligned} \quad (20)$$

since $\bar{X}_i^{(n_0)} < \bar{X}_b^{(n_0)}$ implies $\pi_i \in \mathcal{C}_{\text{aft}}$, $\bar{X}_i^{(n_0)} > \bar{X}_b^{(n_0)}$ implies $\pi_i \in \mathcal{C}_{\text{bef}}$, and $\bar{X}_i^{(n_0)} = \bar{X}_b^{(n_0)}$ has zero probability.

To prove $\Pr(E_i) \leq \alpha/(k-1)$, we introduce two configurations, π_i^* and π_b^* , that are obtained by exchanging the variances of π_i and π_b . That is, π_i^* has mean μ_i and variance σ_b^2 , and π_b^* has mean μ_b and variance σ_i^2 . We will prove that

$$\begin{aligned} &\Pr(E_i \mid \pi_i \in \mathcal{C}_{\text{bef}} \text{ and } \bar{X}_i^{(n_0)} > \bar{X}_b^{(n_0)}) \\ &= \Pr(E_i^* \mid \pi_i^* \in \mathcal{C}_{\text{aft}}^* \text{ and } \bar{X}_i^{*(n_0)} > \bar{X}_b^{*(n_0)}), \end{aligned} \quad (21)$$

where the right-hand side is the probability that π_i^* becomes a provisional best given that (i) \bar{X}' is the sample mean of π_b^* when TSSD starts sampling from π_i^* in Stage 2 (i.e., $\bar{X}_i^{*(n_0)} > \bar{X}_b^{*(n_0)}$) and that (ii) π_i^* has a larger sample mean than π_b^* in Stage 1 (i.e., $\pi_i^* \in \mathcal{C}_{\text{aft}}^*$). Let

$$D_b(r) = \sum_{\ell=1}^r X_{b,\ell} - \frac{r}{N_i} \sum_{\ell=1}^{N_i} X_{i,\ell}$$

and

$$D_i(r) = - \left(\sum_{\ell=1}^r X_{i,\ell}^* - \frac{r}{N_b^*} \sum_{\ell=1}^{N_b^*} X_{b,\ell}^* \right),$$

where $X_{j,\ell}^*$ is the ℓ -th sample from π_j^* for $j = i, b$, and N_b^* is the upper bound, calculated in TSSD, of the number of samples collected from π_b^* . Observe that the left-hand side of Equation (21) is the probability that $D_b(r)$ first crosses the lower edge of the continuation region shown in Figure 2 when S_i^2 is replaced by the sample variance of π_b given that $\bar{X}_i^{(n_0)} > \bar{X}_b^{(n_0)}$. Similarly, the right-hand side is the probability that $D_i(r)$ first crosses the lower edge when S_i^2 is replaced by the sample variance of π_i^* given that $\bar{X}_i^{*(n_0)} > \bar{X}_b^{*(n_0)}$. Therefore Equation (21) holds if $D_b(r)$ and $D_i(r)$ are stochastically equivalent given that $\bar{X}_i^{(n_0)} < \bar{X}_b^{(n_0)}$ and $\bar{X}_i^{*(n_0)} > \bar{X}_b^{*(n_0)}$. Let $X_{j,\ell} = \mu_j + \sigma_j \Delta_{j,\ell}$ for $j = i, b$, $X_{i,\ell}^* = \mu_i + \sigma_b \Delta_{i,\ell}^*$, and $X_{b,\ell}^* = \mu_b + \sigma_i \Delta_{b,\ell}^*$, where $\Delta_{j,\ell}$ and $\Delta_{j,\ell}^*$ are i.i.d. standard normal random variables for $j = i, b$ and $1 \leq \ell \leq n_0$. We now have

$$D_b(r) = \sigma_b \sum_{\ell=1}^r \Delta_{b,\ell} + r \left((\mu_b - \mu_i) - \frac{\sigma_i}{N_i} \sum_{\ell=1}^{N_i} \Delta_{i,\ell} \right)$$

and

$$D_i(r) = -\sigma_b \sum_{\ell=1}^r \Delta_{i,\ell}^* + r \left((\mu_b - \mu_i) + \frac{\sigma_i}{N_b^*} \sum_{\ell=1}^{N_b^*} \Delta_{b,\ell}^* \right).$$

Equation (1) implies that N_b^* is stochastically equivalent to N_i given that $\bar{X}_i^{(n_0)} < \bar{X}_b^{(n_0)}$ and $\bar{X}_i^{*(n_0)} > \bar{X}_b^{*(n_0)}$, since π_b^* and π_i have the same variance, and the sample variance is independent of the sample mean for normal random variables (see Section 4.9 of [Grimmett and Stirzaker 2001]). Now we consider a coupling where $\Delta_{b,\ell} = -\Delta_{i,\ell}^*$ and $\Delta_{i,\ell} = -\Delta_{b,\ell}^*$. Under this coupling,

$$\begin{aligned} \bar{X}_i^{(n_0)} < \bar{X}_b^{(n_0)} &\Leftrightarrow n_0 \mu_i + \sigma_i \sum_{\ell=1}^{n_0} \Delta_{i,\ell} < n_0 \mu_b + \sigma_b \sum_{\ell=1}^{n_0} \Delta_{b,\ell} \\ &\Leftrightarrow n_0 \mu_i + \sigma_b \sum_{\ell=1}^{n_0} \Delta_{i,\ell}^* < n_0 \mu_b + \sigma_i \sum_{\ell=1}^{n_0} \Delta_{i,\ell}^* \\ &\Leftrightarrow \bar{X}_i^{*(n_0)} < \bar{X}_b^{*(n_0)} \end{aligned}$$

and

$$D_i(r) = \sigma_b \sum_{\ell=1}^r \Delta_{b,\ell} + r \left((\mu_b - \mu_i) - \frac{\sigma_i}{N_b^*} \sum_{\ell=1}^{N_b^*} \Delta_{i,\ell} \right).$$

Since $D_i(r)$ and $D_b(r)$ are stochastically equivalent and the conditions, $\bar{X}_i^{(n_0)} < \bar{X}_b^{(n_0)}$ and $\bar{X}_i^{*(n_0)} > \bar{X}_b^{*(n_0)}$, are stochastically equivalent, $D_b(r)$ and $D_i(r)$ are stochastically equivalent given these conditions. This implies Equation (21).

Equations (20) and (21) imply

$$\begin{aligned} \Pr(E_i) &= \Pr(E_i \mid \pi_i \in \mathcal{C}_{\text{aft}} \text{ and } \bar{X}_i^{(n_0)} < \bar{X}_b^{(n_0)}) \Pr(\bar{X}_i^{(n_0)} < \bar{X}_b^{(n_0)}) \\ &\quad + \Pr(E_i^* \mid \pi_i^* \in \mathcal{C}_{\text{aft}}^* \text{ and } \bar{X}_i^{*(n_0)} > \bar{X}_b^{*(n_0)}) \Pr(\bar{X}_i^{*(n_0)} > \bar{X}_b^{*(n_0)}). \end{aligned} \quad (22)$$

By symmetry,

$$\begin{aligned} \Pr(E_i) &= \Pr(E_i \mid \pi_i \in \mathcal{C}_{\text{aft}} \text{ and } \bar{X}_i^{(n_0)} > \bar{X}_b^{(n_0)}) \Pr(\bar{X}_i^{(n_0)} > \bar{X}_b^{(n_0)}) \\ &\quad + \Pr(E_i^* \mid \pi_i^* \in \mathcal{C}_{\text{aft}}^* \text{ and } \bar{X}_i^{*(n_0)} < \bar{X}_b^{*(n_0)}) \Pr(\bar{X}_i^{*(n_0)} < \bar{X}_b^{*(n_0)}). \end{aligned} \quad (23)$$

By taking the average of Equations (22) and (23), we have

$$\begin{aligned} \Pr(E_i) &= \frac{1}{2} \Pr(E_i \mid \pi_i \in \mathcal{C}_{\text{aft}}) + \frac{1}{2} \Pr(E_i^* \mid \pi_i^* \in \mathcal{C}_{\text{aft}}^*) \\ &\leq \frac{\alpha}{k-1}, \end{aligned}$$

where the last inequality follows from Section A.1. This completes the proof of Theorem 2.1 for the case in which $\mathcal{B} = \{\pi_b\}$.

B. PROOF OF THEOREM 2.1 FOR MULTIPLE BEST CONFIGURATIONS

In this section, we discuss the case in which $|\mathcal{B}| \geq 2$. Unfortunately, $\Pr(\pi \in \mathcal{B}) \geq 1 - \alpha$ for $|\mathcal{B}| \geq 2$ does not follow immediately from the arguments in Section A. This is because $\pi_i \in \mathcal{W}$ can become a provisional best while π_b has been eliminated, if $\pi_j \in \mathcal{B}$ eliminates π_b and π_i eliminates π_j . Since $\pi_j \in \mathcal{B}$ can eliminate π_b with probability $> \alpha/(k-1)$ and $\pi_i \in \mathcal{W}$ can eliminate $\pi_j \in \mathcal{B}$ with probability $> \alpha/(k-1)$, it is unclear whether TSSD selects $\pi \in \mathcal{B}$ with probability $\geq 1 - \alpha$. The proof that HN makes a correct selection with a pre-specified high probability when $|\mathcal{B}| \geq 2$ does not appear to formally take into account the such cases [Hong and Nelson 2005].

We will show that TSSD selects $\pi \in \mathcal{B}$ with probability $\geq 1 - \alpha$ for a wide range of parameter settings. Alternatively, TSSD can be modified so that $\pi \in \mathcal{B}$ is guaranteed to be selected with probability $\geq 1 - \alpha$ for any parameter settings. We will show, however, that the modified TSSD is often equivalent to the original TSSD.

We first consider the case when no configurations in \mathcal{B} are sampled before π_b in Stage 2, i.e., $\mathcal{C}_{\text{bef}} \subset \mathcal{W}$. In this case, the argument in Section A implies that π_b becomes a provisional best with probability $\geq 1 - |\mathcal{C}_{\text{bef}}|\alpha/(k-1)$, and that any $\pi_i \in \mathcal{C}_{\text{aft}} \cap \mathcal{W}$ becomes a provisional best with probability $\leq \alpha/(k-1)$, since \bar{X}' can only be updated to a larger value. Therefore, TSSD selects $\pi \in \mathcal{B}$ with probability $\geq 1 - |\mathcal{W}|\alpha/(k-1)$, which is greater than $1 - \alpha$.

Now suppose that a $\pi_j \in \mathcal{B} \setminus \{\pi_b\}$ is sampled before π_b in Stage 2. The argument in Section A implies that a $\pi_w \in \mathcal{W}$ is selected by TSSD with probability at most $\alpha/(k-1)$ if the π_w is sampled before π_j in Stage 2, since the π_w is more likely to be eliminated by π_j and π_b than by π_b alone. However, the argument in Section A does not immediately imply that a $\pi_w \in \mathcal{W}$ is selected by TSSD with probability $\leq \alpha/(k-1)$ if a $\pi_j \in \mathcal{B} \setminus \{\pi_b\}$ is sampled before the π_w in Stage 2. Suppose that a $\pi_j \in \mathcal{B} \setminus \{\pi_b\}$ becomes a provisional best and TSSD updates $\bar{X}' = \bar{X}_j^{(N_j)}$.

Now a π_w becomes a provisional best if $r(\bar{X}_w^{(r)} - \bar{X}')$ exits the continuation region from the upper edge, which can occur with probability $\geq \alpha/(k-1)$, and \bar{X}' is updated as $\bar{X}' = \max\{\bar{X}', \bar{X}_w^{(N_w)}\}$. Then π_b becomes a provisional best only if $r(\bar{X}_b^{(r)} - \bar{X}')$ exits the continuation region from the lower edge. In Section A, we showed that π_b becomes a provisional best with probability $\geq 1 - \alpha/(k-1)$ if $\bar{X}' = \bar{X}_w^{(N_w)}$. Since $\bar{X}' = \max\{\bar{X}', \bar{X}_w^{(N_w)}\} \geq \bar{X}_w^{(N_w)}$, the probability that π_b becomes a provisional best decreases from when $\bar{X}' = \bar{X}_w^{(N_w)}$. Therefore, the error probability, $\Pr(\text{ICS}_w)$, increases. Observe however that a $\pi \in \mathcal{B}$ is selected with probability $\geq 1 - \sum_{\pi_w \in \mathcal{W}} \Pr(\text{ICS}_w)$. Therefore, even though each $\Pr(\text{ICS}_w)$ may become larger, the sum of $\Pr(\text{ICS}_w)$ may be smaller when $|\mathcal{B}| \geq 2$ than when $|\mathcal{B}| = 1$.

We will now formalize the above argument. Let $V(\pi_i, x)$ be the event that there exists an r such that $\bar{X}_i^{(r)} - x < -W_i(r)$ and $r < s$ for any s such that $\bar{X}_i^{(s)} - x > W_i(s)$ (i.e., $\bar{X}_i^{(r)} - x$ exits the continuation region from the lower edge, $-W_i(r)$). Let $\Lambda(\pi_i, x)$ be the event that there exists an r such that $\bar{X}_i^{(r)} - x > W_i(r)$ and $r < s$ for any s such that $\bar{X}_i^{(s)} - x < -W_i(s)$ (i.e., $\bar{X}_i^{(r)} - x$ exits the continuation region from the upper edge, $W_i(r)$).

Let $P_w^{\text{bef}}(x)$ be the probability that TSSD selects $\pi_w \in \mathcal{W} \cap \mathcal{C}_{\text{bef}}$ as the best configuration when $\bar{X}' = x$ immediately before sampling from π_w in Stage 2. Since TSSD selects the π_w only if the π_w becomes a provisional best and π_w is still the provisional best after samples from π_b are collected in Stage 2,

$$P_w^{\text{bef}}(x) \leq \Pr\left(\Lambda(\pi_w, x) \text{ and } V(\pi_b, \max\{x, \bar{X}_w^{(N_w)}\})\right).$$

The value of \bar{X}' is at least $\max\{x, \bar{X}_w^{(N_w)}\}$ when we start sampling from π_b in Stage 2 given that π_w becomes a provisional best. By conditioning on whether or not $x \geq \bar{X}_w^{(N_w)}$, we have

$$\begin{aligned} P_w^{\text{bef}}(x) &\leq \Pr\left(\Lambda(\pi_w, x) \text{ and } x \geq \bar{X}_w^{(N_w)} \text{ and } V(\pi_b, x)\right) \\ &\quad + \Pr\left(\Lambda(\pi_w, x) \text{ and } x < \bar{X}_w^{(N_w)} \text{ and } V(\pi_b, \bar{X}_w^{(N_w)})\right) \\ &\leq \Pr\left(\Lambda(\pi_w, x) \text{ and } V(\pi_b, x)\right) + \Pr\left(V(\pi_b, \bar{X}_w^{(N_w)})\right) \\ &= \Pr\left(\Lambda(\pi_w, x)\right) \Pr\left(V(\pi_b, x)\right) + \Pr\left(V(\pi_b, \bar{X}_w^{(N_w)})\right), \end{aligned} \quad (24)$$

where the last equality follows from the independence of $\Lambda(\pi_w, x)$ and $V(\pi_b, x)$.

Let $P_w^{\text{aft}}(x)$ be the probability that TSSD selects $\pi_w \in \mathcal{W} \cap \mathcal{C}_{\text{aft}}$ as the best configuration when $\bar{X}' = x$ immediately before sampling from π_b in Stage 2. Note that TSSD selects the π_w only if the π_w becomes a provisional best. By conditioning on whether or not π_b becomes a provisional best,

$$P_w^{\text{aft}}(x) \leq \Pr\left(V(\pi_b, x) \text{ and } \Lambda(\pi_w, x)\right) + \Pr\left(\Lambda(\pi_b, x) \text{ and } \Lambda(\pi_w, \max\{x, \bar{X}_b^{(N_b)}\})\right),$$

since the value of \bar{X}' is at least x if π_b does not become a provisional best and is at least $\max\{x, \bar{X}_b^{(N_b)}\}$ if π_b becomes a provisional best. Therefore, as with the

previous argument, we have

$$P_w^{\text{aft}}(x) \leq \Pr(V(\pi_b, x)) \Pr(\Lambda(\pi_w, x)) + \Pr\left(\Lambda(\pi_w, \max\{x, \bar{X}_b^{(N_b)}\})\right). \quad (25)$$

Let P_w be the probability that TSSD selects a $\pi_w \in \mathcal{W}$ as the best configuration. The arguments in Section A can be used to show that the last terms of Equation (24) and Equation (25) are at most $\alpha/(k-1)$. Therefore, it follows that $P_w \leq \max_x P_w(x)$, where

$$P_w(x) = \Pr(V(\pi_b, x)) \Pr(\Lambda(\pi_w, x)) + \frac{\alpha}{k-1}.$$

Observe that $\Pr(V(\pi_b, x)) < 1/2$ if $x > \mu_b$ and that $\Pr(\Lambda(\pi_w, x)) < 1/2$ if $x < \mu_w$. Since Lemma C.4 in Section C can be proved in a way similar to Inequality (18), we have

$$\Pr(V(\pi_b, x)) \leq \begin{cases} \mathbb{E}\left[\left(1 + \exp\left(\sqrt{\frac{h^2 Q}{n_0-1}} Z + \frac{h^2 Q}{n_0-1} \frac{\mu_b - x}{\delta}\right)\right)^{-1}\right] & \text{if } x \leq \mu_b \\ \frac{1}{2} & \text{if } x > \mu_b \end{cases} \quad (26)$$

and

$$\Pr(\Lambda(\pi_w, x)) \leq \begin{cases} \mathbb{E}\left[\left(1 + \exp\left(\sqrt{\frac{h^2 Q}{n_0-1}} Z + \frac{h^2 Q}{n_0-1} \frac{x - \mu_w}{\delta}\right)\right)^{-1}\right] & \text{if } x \geq \mu_w \\ \frac{1}{2} & \text{if } x < \mu_w, \end{cases} \quad (27)$$

where Q is a χ^2 random variable with $n_0 - 1$ degrees of freedom, Z is a standard normal random variable, and Q and Z are independent. Note that Lemma C.4 provides a bound on the probability that a random walk with a deterministic drift, x , crosses a lower edge before crossing an upper edge, while Lemma C.1 used in Section A provides a corresponding bound for a random walk with a stochastic drift.

Let $P_{\mathcal{W}}$ be the probability that TSSD selects one of the configurations in \mathcal{W} . Observe that $P_{\mathcal{W}} \leq \sum_{\pi_w \in \mathcal{W}} P_w$. When $|\mathcal{B}| \geq 3$, $P_{\mathcal{W}}$ can be shown to be smaller than α for a wide range of parameter settings. When $|\mathcal{B}| = 2$, however, a tighter analysis is often needed. Let $\mathcal{B} = \{\pi_b, \pi_j\}$ so that $\mu_b - \delta < \mu_j \leq \mu_b$. Since we reorder the configurations in decreasing order of the sample means at the end of Stage 1, π_b is sampled before π_j in Stage 2 with probability $\geq 1/2$. If π_b is sampled before π_j , we have seen above that a $\pi_w \in \mathcal{W}$ is selected by TSSD with probability $\leq \alpha/(k-1)$. When $|\mathcal{B}| = 2$, we thus have

$$P_{\mathcal{W}} \leq \frac{1}{2}(k-2) \frac{\alpha}{k-1} + \frac{1}{2} \sum_{\pi_w \in \mathcal{W}} P_w.$$

Therefore, for any $|\mathcal{B}|$, we have

$$P_{\mathcal{W}} \leq \max \left\{ (k-3)P_w, \frac{1}{2}(k-2) \frac{\alpha}{k-1} + \frac{1}{2}(k-2)P_w \right\}, \quad (28)$$

where $P_w \leq \max_x P_w(x)$ and $P_w(x)$ is bounded from above via Inequalities (26)-(27).

Figure 11 shows an upper bound of $P_{\mathcal{W}}$. Specifically, the figure plots, as a function of x/δ , the right-hand side of Inequality (28) with P_w replaced by the upper bound




Fig. 11. An upper bound of $P_{\mathcal{W}}$ from Inequality (28). We set $\alpha = 0.05$.

of $P_w(x)$ given via Inequalities (26)-(27). We set $k = 10$, $n_0 = 10$, and $\alpha = 0.05$. In these settings, the confidence parameter is $h = 4.69$. The figures suggest that $P_{\mathcal{W}} \leq \alpha$ for these parameter settings. We omit the figure here, but we found that $P_{\mathcal{W}} \leq \alpha$ when $k = 10$, $n_0 = 10$, and $\alpha = 0.1$. Of course, $P_{\mathcal{W}} \leq \alpha$ does not hold for all parameter settings, particularly when k is large. Note, however, that the above analysis is not tight, and $P_{\mathcal{W}}$ may be much smaller than the upper bound derived above. Also, one may set h larger so that $P_{\mathcal{W}} \leq \alpha$ is guaranteed for a given parameter setting.

C. TECHNICAL LEMMAS AND PROOFS

In this section, we present technical lemmas used in Section A, and prove Proposition 2.3.

LEMMA C.1. Let a , c , $L(r)$, and $U(r)$ be as defined in Lemma 2.2. Let X_1, X_2, \dots be standard normal random variables and Δ be a normal random variable with mean $\mu \geq 0$ and standard deviation $\sigma \geq 0$. We assume that Δ and the X_ℓ 's are independent. Let $S_\Delta(r) = \sum_{\ell=1}^r X_\ell + r \Delta$, and E be the event that $S_\Delta(r)$ crosses the lower edge, $L(r)$, before crossing the upper edge, $U(s)$. Formally, E is the event that there exists an r such that $S_\Delta(r) < L(r)$ and $r < s$ for any s such that $S_\Delta(s) > U(s)$. Then

$$\Pr(E) \leq \mathbf{E} \left[\frac{1}{1 + \exp(2(\sqrt{ac}Z + c\Delta))} \right]$$

where Z is a standard normal random variable that is independent of Δ .

PROOF. Let $S_x(r) = \sum_{\ell=1}^r X_\ell + r x$, and let E_x be the event that $S_x(r)$ crosses a lower edge, $L(r)$, before crossing an upper edge, $U(r)$. Formally, E_x is the event that there exists an r such that $S_x(r) < L(r)$ and $r < s$ for any s such that

$S_x(s) > U(s)$. Lemma 2.2 can be used to show that (see p. 170 of [Anderson 1960]) there exists a function $\epsilon(x)$ satisfying $\epsilon(x) < 0$ for $x > 0$, $\epsilon(x) > 0$ for $x < 0$, and $\epsilon(0) = 0$ such that

$$\Pr(E_x) = \mathbb{E} \left[\frac{1}{1 + \exp(2(\sqrt{ac}Z + cx))} \right] + \epsilon(x). \quad (29)$$

Observe that

$$\begin{aligned} \Pr(E) &= \int_{-\infty}^{\infty} \Pr(E_x) \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx \\ &= \mathbb{E} \left[\frac{1}{1 + \exp(2(\sqrt{ac}Z + c\Delta))} \right] + \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} \epsilon(x) e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx. \end{aligned} \quad (30)$$

Notice that $S_x(r)$ can be coupled with $S_{-x}(r)$ so that $S_x(r)$ crosses $L(r)$ whenever $S_{-x}(r)$ crosses $U(r)$ and vice versa. Then it follows that

$$\begin{aligned} \Pr(E_{-x}) &= \Pr(S_{-x}(r) \text{ crosses } L(r) \text{ before crossing } U(r)) \\ &= \Pr(S_x(r) \text{ crosses } U(r) \text{ before crossing } L(r)) \\ &= 1 - \Pr(E_x). \end{aligned}$$

Since $\Pr(E_x) + \Pr(E_{-x}) = 1$, Equation (29) can be used to show that

$$\mathbb{E} \left[\frac{1}{1 + \exp(2(\sqrt{ac}Z + cx))} \right] + \mathbb{E} \left[\frac{1}{1 + \exp(2(\sqrt{ac}Z - cx))} \right] + \epsilon(x) + \epsilon(-x) = 1.$$

The first term is the probability that the Brownian motion with drift x crosses $L(r)$ before crossing $U(r)$. The second term is the probability that the Brownian motion with drift $-x$ crosses $L(r)$ before crossing $U(r)$, which is equivalent to the probability that the Brownian motion with drift x crosses $U(r)$ before crossing $L(r)$. Therefore, the sum of the first two terms is one. Thus, $\epsilon(x) = -\epsilon(-x)$, which implies that

$$\begin{aligned} \int_{-\infty}^{\infty} \epsilon(x) e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx &= \int_{-\infty}^0 \epsilon(x) e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx - \int_0^{-\infty} \epsilon(-x) e^{-\frac{(-x-\mu)^2}{2\sigma^2}} dx \\ &= \int_{-\infty}^0 \epsilon(x) \left(e^{-\frac{(x-\mu)^2}{2\sigma^2}} - e^{-\frac{(x+\mu)^2}{2\sigma^2}} \right) dx \\ &\leq 0, \end{aligned} \quad (31)$$

where the last inequality follows from the fact that $\epsilon(x) > 0$ for $x < 0$. Equations (30) and (31) prove the lemma. \square

LEMMA C.2. *Let Z be a standard normal random variable. Then*

$$\mathbb{E} \left[\frac{1}{1 + e^{\sigma Z + \mu}} \right]$$

is a nondecreasing function of σ for $\sigma \geq 0$ and $\mu \geq 0$.

PROOF. Note that $Y = \sigma Z + \mu$ is a normal random variable, $N(\mu, \sigma)$. Thus,

$$\frac{\partial}{\partial \sigma} \mathbb{E} \left[\frac{1}{1 + e^{\sigma Z + \mu}} \right] = \frac{\partial}{\partial \sigma} \int_{-\infty}^{\infty} \frac{1}{1 + e^y} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y-\mu)^2}{2\sigma^2}} dy.$$

The dominated convergence theorem can be used to show that the derivative and the integral can be exchanged. Hence,

$$\begin{aligned}
 \frac{\partial}{\partial \sigma} \mathbb{E} \left[\frac{1}{1 + e^{\sigma Z + \mu}} \right] &= \frac{1}{\sigma} \int_{-\infty}^{\infty} \frac{1}{1 + e^y} \frac{1}{\sqrt{2\pi\sigma^2}} \left(\left(\frac{y - \mu}{\sigma} \right)^2 - 1 \right) e^{-\frac{(y - \mu)^2}{2\sigma^2}} dy \\
 &= \frac{1}{\sigma} \mathbb{E} \left[\frac{1}{1 + e^Y} \left(\left(\frac{Y - \mu}{\sigma} \right)^2 - 1 \right) \right] \\
 &= \frac{1}{\sigma} \mathbb{E} \left[\frac{Z^2 - 1}{1 + e^{\sigma Z + \mu}} \right]. \tag{32}
 \end{aligned}$$

It remains to prove that the right-hand side, $\psi(\mu)$, of Equation (32) is nonnegative for $\mu \geq 0$ and $\sigma \geq 0$. Below, we will show $\psi(0) = 0$ and $\psi(\mu) \geq 0$ for any $\mu > 0$.

First, consider $\psi(0)$. By symmetry of the standard normal random variable,

$$\psi(0) = \frac{1}{\sigma} \mathbb{E} \left[\frac{(-Z)^2 - 1}{1 + e^{-\sigma Z}} \right] = \frac{1}{\sigma} \mathbb{E} \left[\frac{(Z^2 - 1)e^{\sigma Z}}{e^{\sigma Z} + 1} \right].$$

Summing $\psi(0)$ and the last expression, we obtain

$$\begin{aligned}
 2\psi(0) &= \frac{1}{\sigma} \mathbb{E} \left[\frac{Z^2 - 1}{1 + e^{\sigma Z}} \right] + \frac{1}{\sigma} \mathbb{E} \left[\frac{(Z^2 - 1)e^{\sigma Z}}{e^{\sigma Z} + 1} \right] \\
 &= \frac{\mathbb{E} [Z^2] - 1}{\sigma} \\
 &= 0. \tag{33}
 \end{aligned}$$

Next, consider $\psi(\mu)$ for $\mu > 0$.

$$\begin{aligned}
 \psi(\mu) &= \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} \frac{z^2 - 1}{1 + e^{\sigma z + \mu}} e^{-\frac{z^2}{2}} dz \\
 &= \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^{\infty} (z^2 - 1) e^{-\frac{z^2}{2}} \xi_{\mu}(z) dz
 \end{aligned}$$

where we define

$$\xi_{\mu}(z) \equiv \frac{1}{1 + e^{\sigma z + \mu}} + \frac{1}{1 + e^{-\sigma z + \mu}}.$$

Note that $\xi_{\mu}(z)$ is a nondecreasing function of z for $z \geq 0$, since

$$\xi_{\mu}'(z) = \frac{\sigma e^{\mu} (e^{\sigma z} - e^{-\sigma z}) (e^{2\mu} - 1)}{(1 + e^{\sigma z + \mu})^2 (1 + e^{-\sigma z + \mu})^2} \geq 0$$

for $\mu \geq 0$, $\sigma \geq 0$, and $z \geq 0$. Hence, $\xi_{\mu}(z) \leq \xi_{\mu}(1)$ for $z \leq 1$, and $\xi_{\mu}(z) \geq \xi_{\mu}(1)$ for $z \geq 1$. Since $(z^2 - 1)e^{-\frac{z^2}{2}} \leq 0$ for $z \leq 1$ and $(z^2 - 1)e^{-\frac{z^2}{2}} \geq 0$ for $z \geq 1$,

$$\begin{aligned}
 \psi(\mu) &\geq \frac{1}{\sqrt{2\pi\sigma^2}} \left(\int_0^1 (z^2 - 1) e^{-\frac{z^2}{2}} \xi_{\mu}(1) dz + \int_1^{\infty} (z^2 - 1) e^{-\frac{z^2}{2}} \xi_{\mu}(1) dz \right) \\
 &= \frac{\xi_{\mu}(1)}{\sqrt{2\pi\sigma^2}} \int_0^{\infty} (z^2 - 1) e^{-\frac{z^2}{2}} dz \\
 &= \xi_{\mu}(1) \psi(0),
 \end{aligned}$$

where the last inequality follows from Equation (33). Since $\psi(0) = 0$, this completes the proof of the lemma. \square

LEMMA C.3. *Let N be a random variable with support on $\{n_0, n_0+1, \dots\}$, where n_0 is a natural number. Let X_1, X_2, \dots be i.i.d. normal random variables with mean μ and σ such that $\sum_{\ell=1}^r X_\ell$ and N are independent for any $r \geq n_0$. Then $Z^{(N)} = \left(\sum_{\ell=1}^N X_\ell\right)/N - \mu / \sqrt{\sigma^2/N}$ is a standard normal random variable. Also, $Z^{(N)}$ and N are independent.*

PROOF. Let $\Phi(\cdot)$ be the standard normal distribution function. Observe that $\Pr(Z^{(r)} \leq z) = \Phi(z)$ for any constant r . Also, $Z^{(r)}$ and N are independent for any constant $r \geq n_0$, since $\sum_{\ell=1}^r X_\ell$ and N are independent for $r \geq n_0$. Hence,

$$\begin{aligned} \Pr(Z^{(N)} \leq z) &= \sum_{r \geq n_0} \Pr(Z^{(r)} \leq z \mid N = r) \Pr(N = r) \\ &= \sum_{r \geq n_0} \Pr(Z^{(r)} \leq z) \Pr(N = r) \\ &= \sum_{r \geq n_0} \Phi(z) \Pr(N = r) \\ &= \Phi(z) \end{aligned}$$

Thus, $Z^{(N)}$ has a standard normal distribution. Also, $Z^{(N)}$ and N are independent, since $Z^{(N)}$ is a standard normal random variable for any given $N = r \geq n_0$. \square

LEMMA C.4. *Let $\Lambda(\pi_i, x)$ and $V(\pi_i, x)$ be as defined in Section B. If $x \leq \mu_b$, then we have*

$$\Pr(V(\pi_b, x)) \leq \mathbf{E} \left[\frac{1}{1 + \exp \left(\sqrt{\frac{h^2 Q}{n_0 - 1}} Z + \frac{h^2 Q}{n_0 - 1} \frac{\mu_b - x}{\delta} \right)} \right], \quad (34)$$

where Q is a χ^2 random variable with $n_0 - 1$ degrees of freedom, Z is a standard normal random variable, and Q and Z are independent. Also, if $\mu_w \leq x$, then we have

$$\Pr(\Lambda(\pi_w, x)) \leq \mathbf{E} \left[\frac{1}{1 + \exp \left(\sqrt{\frac{h^2 Q}{n_0 - 1}} Z + \frac{h^2 Q}{n_0 - 1} \frac{x - \mu_w}{\delta} \right)} \right]. \quad (35)$$

PROOF. We prove only Inequality (34), since Inequality (35) follows from the symmetry. Recall that $V(\pi_b, x)$ is the event that $\bar{X}_b^{(r)} - x$ crosses a lower edge, $-W_b(r)$, before crossing an upper edge, $W_b(r)$, so that

$$\begin{aligned} &\Pr(V(\pi_b, x)) \\ &= \Pr \left(\exists r \leq N_b \text{ s.t. } \bar{X}_b^{(r)} < x - W_b(r) \text{ and } r < s \forall s \text{ s.t. } \bar{X}_b^{(s)} > x + W_b(s) \right), \end{aligned}$$

where we recall that

$$N_b = \max \left\{ n_0, \left\lceil \frac{h^2 S_b^2}{\delta^2} \right\rceil \right\}$$

$$\bar{X}_b^{(r)} = \sum_{\ell=1}^r X_{b,\ell}$$

$$W_b(r) = \max \left\{ 0, \frac{\delta}{2r} \left(\frac{h^2 S_b^2}{\delta^2} - r \right) \right\}.$$

Let $Z_\ell = (X_{b,\ell} - \mu_b)/\sigma_b$. Then

$$\Pr(V(\pi_b, x)) = \Pr \left(\exists r \leq N_b \text{ s.t. } \frac{\sigma_b}{r} \sum_{\ell=1}^r Z_\ell < x - W_b(r) - \mu_b \right. \\ \left. \text{and } r < s \forall s \text{ s.t. } \frac{\sigma_b}{r} \sum_{\ell=1}^r Z_\ell > x + W_b(r) - \mu_b \right).$$

By Equation (3), we have

$$\Pr(V(\pi_b, x)) \\ = \Pr \left(\exists r \leq N_b \text{ s.t. } \sum_{\ell=1}^r Z_\ell + \frac{r}{\sigma_b} (\mu_b - x) < -\max \left\{ 0, \frac{h^2 S_b^2}{2\delta\sigma_b} - \frac{\delta}{2\sigma_b} r \right\} \right. \\ \left. \text{and } r < s \forall s \text{ s.t. } \sum_{\ell=1}^r Z_\ell + \frac{r}{\sigma_b} (\mu_b - x) > \max \left\{ 0, \frac{h^2 S_b^2}{2\delta\sigma_b} - \frac{\delta}{2\sigma_b} r \right\} \right).$$

Given S_b^2 , $a = \delta/(2\sigma_b)$ and $c = h^2 S_b^2/(2\delta\sigma_b)$ are constants. Hence Lemma 2.2 can be used to show that (see p. 170 of [Anderson 1960])

$$\Pr(V(\pi_b, x)) \leq \mathbb{E} \left[\mathbb{E} \left[\frac{1}{1 + \exp \left(2 \left(\sqrt{\frac{h^2 S_b^2}{2\delta\sigma_b} \frac{\delta}{2\sigma_b}} Z + \frac{h^2 S_b^2}{2\delta\sigma_b} \frac{\mu_b - x}{\sigma_b} \right) \right)} \middle| S_b^2 \right] \right].$$

Since $Q = S_b^2/(n_0 - 1)$ is a χ^2 random variable with $n_0 - 1$ degrees of freedom, this implies Inequality (34). \square

PROOF OF PROPOSITION 2.3. Let $\Phi(\cdot)$ be the standard normal distribution function. We use Inequality (4.65) from [Anderson 1960]:

$$\left| \Phi(z) - \frac{1}{1 + e^{-\kappa z}} \right| < 0.01, \quad (36)$$

for $-\infty < z < \infty$. By Inequality (36),

$$\mathbb{E} \left[\frac{1}{1 + e^{\sigma Z + \mu}} \right] \leq \int_{-\infty}^{\infty} \left(\Phi \left(-\frac{\sigma z + \mu}{\kappa} \right) + 0.01 \right) \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz \\ = \int_{-\infty}^{\infty} \Phi \left(-\frac{\sigma z + \mu}{\kappa} \right) \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz + 0.01 \\ = \Pr \left(Z' \leq -\frac{\sigma Z'' + \mu}{\kappa} \right) + 0.01, \quad (37)$$

where Z' and Z'' are independent standard normal random variables. Let $Z''' = (Z' + \sigma Z''/\kappa)/\sqrt{1 + \sigma^2/\kappa^2}$. Then Z''' is a standard normal random variable. By

Inequality (37),

$$\mathbb{E} \left[\frac{1}{1 + e^{\sigma Z + \mu}} \right] \leq \Pr \left(Z''' \leq -\frac{\mu}{\sqrt{\kappa^2 + \sigma^2}} \right) + 0.01,$$

which proves the upper bound of Inequality (6). The lower bound can be proved similarly.

Also, by a property of the normal distribution function, $g(h)$ is a decreasing function of h . \square