Christian Schmidt

Evolutionary Computation in Stochastic Environments

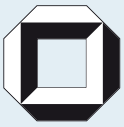

universitätsverlag karlsruhe

Christian Schmidt

Evolutionary Computation in Stochastic Environments

Evolutionary Computation in Stochastic Environments

von
Christian Schmidt

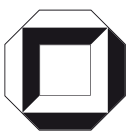

universitätsverlag karlsruhe

Dissertation, genehmigt von der Fakultät für Wirtschaftswissenschaften der
Universität Fridericiana zu Karlsruhe, 2006

Referenten: Prof. Dr. H. Schmeck
Prof. Dr. K.-H. Waldmann
Prof. S. E. Chick, Ph.D.

Impressum

Universitätsverlag Karlsruhe
c/o Universitätsbibliothek
Straße am Forum 2
D-76131 Karlsruhe
www.uvka.de

Universitätsverlag Karlsruhe 2007
Print on Demand

ISBN: 978-3-86644-128-6

meiner Frau

Danksagung

Die vorliegende Arbeit entstand während meiner Tätigkeit als wissenschaftlicher Mitarbeiter des Instituts AIFB der Universität Karlsruhe (TH) im Rahmen des Kooperationsprojektes “Supply Chain Planning” mit der Firma LOCOM Consulting GmbH, Karlsruhe.

Mein Dank gilt meinem Betreuer Professor Schmeck, der mir das Vertrauen, die Freiräume und die organisatorischen und finanziellen Rahmenbedingungen gab, mich in meinem Forschungsthema so zu entfalten, dass ich eine mich erfüllende und faszinierende Promotion anfertigen konnte.

Weiterhin danke ich dem Gründer der Firma LOCOM, Volker Klohr, dass er den Kontakt zum AIFB seit vielen Jahren durch Vorträge, Seminare und Diplomarbeiten pflegt und sich im Rahmen der Kooperation zur Finanzierung meiner Promotionsstelle bereit erklärte. Ich habe große Achtung vor der Bereitschaft sich als kleines mittelständisches Unternehmen in der Forschung in diesem Umfang zu engagieren. Um so mehr freue ich mich, dass LOCOM die Kooperation als einen großen Erfolg betrachtet. Auch für mich war die Erfahrung, Optimierungsprobleme in der Praxis zu lösen, zugleich lehrreich und inspirierend. Sehr wichtig war auch die mir eingeräumte Zeit, um mehr als nur eine “Freizeit”-Promotion anzufertigen und vor allen Dingen die Anteilnahme und Anerkennung, die ich durch Volker Klohr und meine Kollegen (“wir sind Doktor”) erfahren durfte.

Ganz besonders danke ich Jürgen Branke für die intensive Begleitung bei der Promotion, die vielen angeregten Diskussionen und die gemeinsamen Publikationen. Er verstand es stets meinen kreativen Ausschweifungen den nötigen wissenschaftlichen Unterbau zu verpassen und den Blick für das Ziel zu bewahren. Aber am allermeisten danke ich ihm dafür, mich mit dem

"Forscher-Virus" lebenslänglich infiziert zu haben.

A big "Dankeschön" to Steve Chick for the intensive and refreshing research work we have done together, for realtime email correspondence at night and especially for surprising insights into Bayesian statistics. Thanks for serving as a referent of my thesis.

Meinen Kollegen an der Uni danke ich nicht nur für das Nicht-verdurstenlassen meiner Pflanzen, sondern auch für die Bereitschaft, meine Leidenschaft für türkisches Essen und Abneigung gegen Mensaessen zu akzeptieren. Ich habe die Atmosphäre am Institut und das Zusammensein immer außerordentlich genossen.

Ich danke meinen Eltern, dass sie die Geduld für meine lange Ausbildung aufgebracht haben, dass sie mir das Gefühl geben, auf mich stolz zu sein, mich bestärkt haben in dem Gedanken zu promovieren und dass sie die Eltern sind, die sie sind.

Mein größter Dank gilt der Frau, die ich liebe und der ich diese Arbeit gewidmet habe. Sie hat mich in der stressigen Phase der Arbeit physisch und psychisch am Leben gehalten, mir all die "Trivialitäten", die das Leben bereit hält, abgenommen und dabei gleichzeitig Verständnis und Rücksicht aufgebracht.

Karlsruhe, im Februar 2007 Christian Schmidt

Contents

Chapter 1

Introduction

Evolutionary Algorithms (EA) are iterative optimization heuristics, which are inspired by natural evolution. They have proven to be very successful in many different application areas including logistics, telecommunication and production planning. One of their particular advantages is the wide applicability, as they do not impose strong restrictions like continuity or convexity on the model used for optimization. Normally, EA are applied to deterministic problems, i.e. problems with all parameters and outputs known exactly. In reality, decision makers are often confronted with problems where this is not the case. Uncertainties might arise, if the problem contains parameters that are outcomes of future developments and therefore unpredictable, e.g. weather or customer demand. Further uncertainties arise from parameters or outputs that might be determined exactly, but at costs exceeding the possible gain in value.

In practice, these uncertain parameters are often replaced with their expected values and the problem is solved as a deterministic problem. If the unknown parameter is critical, then the solution might be very sensitive to deviations of the uncertain parameter from its expected value, and therefore the solution might have a bad value when implemented.

Problems with uncertain parameters can be modeled as stochastic problems. The uncertain parameters are replaced by stochastic variables with distributions representing the decision maker's belief in the possible realiza-

tions. The solution for the stochastic problem will be optimized for having the best expected value or utility over all uncertain parameters.

The application of optimization procedures designed for deterministic problems to stochastic problems is not straightforward, because, in general, the value of a solution can not be determined exactly. It can only be estimated by repeated evaluations of different scenarios, where a scenario in this context means a possible realization of the uncertain parameters. For complex problems this evaluation is often done by simulating the behavior of a solution for the given scenario. Simulation is a very powerful tool, as it allows to model complex systems in a natural and realistic way.

Adapting EA for stochastic problems is one promising possibility to combine optimization and stochastic simulation for solving complex and flexible models of real world problems with combinatorially large search spaces. Two specialties of EA further increase the elegance of this combination: Firstly, EA can optimize multiple criteria in one run, so decision makers are not forced to give an a-priori weighting of concurring measures of performance like time and cost. Secondly, the evaluation of solutions can be distributed easily on different nodes of a computer network with modest requirements on communication among the nodes, which allows the use of standard PC components to improve the speed of the optimization process.

In this thesis it is shown how to design EA for solving stochastic optimization problems by combining features of statistical ranking and selection procedures with features of stochastic optimization procedures.

1.1 Scope

The general setting of this thesis is to find the configuration

$$x^* = \arg\max_{x \in X} \mathbf{E}_{\omega}[f(x, \omega)], \tag{1.1}$$

where x denotes a configuration (represented by a vector of input variables) from the search space X, ω a random scenario (represented by a vector of environment variables) and f the performance measure of a given configuration

in a certain scenario. The objective is to find the configuration with the best expected performance over all scenarios. In various contexts configurations are called settings, designs, systems or solutions. This model is equivalent to the formulation of "Simulation Optimization" used by [Fu 2002].

The given formulation indicates that the focus is on finding the best configuration x^* rather than the exact value of the objective function, which conforms to the needs of practitioners. For the search space it is assumed that configurations can be generated easily as opposed to feasibility problems, where considerable effort goes into finding valid configurations. The objective must be estimatable by repeated evaluations of the performance measure for different scenarios, i.e. the sample average is a consistent estimator for the quality of a configuration. If the performance of competing configurations can be distinguished easily by few evaluations, the problem can be solved as a quasi-deterministic problem and the approach presented here provides no additional benefit.

Although no proofs for convergence of EA on stochastic problems, except for few special cases, are known yet, many commercial software packages for simulation make use of EA and related approaches to generate and select good solutions. For simulation optimization the EA's ability to perform "black-box" optimization, i.e. optimization with few assumptions on the problem structure, is advantageous.

1.2 Related Approaches

The model used in the formulation of the stochastic optimization problem above is very general. The approaches to solve Equation 1.1 differ in the importance given to certain aspects of the problem and in the assumptions made on the structure of the performance measure and the search space.

Statistical Decision Theory Statistical decision theory developed by [Wald 1950] is the basis for many ranking and selection procedures. They are introduced in more detail in Chapter 2. For a given set of alternative systems a probability of correct selection is guaranteed approximately or on

average. These approaches are suitable, when the number of alternatives is small and an evaluation of the performance is expensive. Our approach combines EA with ranking and selection procedures in order to allow search spaces of combinatorial size instead of a few alternatives.

Markov Decision Processes Another approach considers dynamic programming and Markov decision processes (see [Kall and Wallace 1994] for an introduction). Here, one searches for optimal actions to execute at discrete points in time. The actions generate random outcome and transform the current state to the next stage probabilistically. The number of actions, states and stages needs to be low, as the general approach is to form a backward recursion to identify optimal actions for each state.

Stochastic Programming In the operations research (OR) community typically problems are solved by Stochastic Programming (see [Birge and Louveaux 1997]). The models used are generally restricted to linear objective functions and constraints, but nonetheless many real-world problems are too complex to be solved. For stochastic optimization problems the optimality of a solution can not be guaranteed any more as opposed to deterministic ones, but only a probability that the solution is near the optimal solution. The number of scenarios (realizations of the random parameters) is usually fixed a-priori.

Stochastic Programming as introduced by [Birge and Louveaux 1997] is the basis for many algorithms in classical OR. The general formulation is given by two nested linear programs, where the second stage problems are solved conditional on a given scenario and linked to the first stage by the expected value, a convex, non-linear function that cannot be calculated exactly.

The Sample Average Approximation (SAA) by [Kleywegt, Shapiro, and de Mello 2001] estimates the expected value of the second-stage problems by approximating it with a fixed number N of scenarios and solving the resulting problem with Benders' Decomposition ([Benders 1962]).

For the SAA to achieve an ϵ-optimal solution with probability at least

$1 - \alpha$, the number of scenarios must be $N \geq \frac{\sigma^2}{\epsilon^2} \log \frac{|Y|}{\alpha}$ with σ^2 being the problem specific variance and $|Y|$ the problem size. That means the deviation ϵ decreases with $O(N^{-0.5})$ or in other words: to double the accuracy the number of scenarios must be quadrupled.

1.3 Motivation

In this section we will motivate why we believe that EA are well suited for solving large scale stochastic optimization problems and which research results might positively influence the efficiency of EA on these problems.

For the problems emphasized in this thesis, it is assumed that the evaluation of a configuration or solution is expensive in terms of time and cost. So to solve problems with large search spaces it is indispensable that the number of scenarios required for measuring a solution's performance is reduced as much as possible without deteriorating the optimization procedure's mechanisms. Using too few scenarios would result in a non-effective procedure while too many would lead to an inefficient procedure.

Several ideas for an efficient allocation of the number of necessary scenarios or samples are known from other areas of research. Combining these approaches gives reason to hope that the number of scenarios needed for optimization can be reduced significantly compared to the usual strategy, where the number of samples is fixed a-priori, based on the variance of the mean estimator. The following enumeration sketches theses ideas in short.

1. **Ordinal Optimization**

   [Ho, Sreenivas, and Vakili 1992] noted that it is easier to compare solutions than to estimate their performance precisely. Actually the probability that the means of two random variables X and Y differ is $P(\bar{X}_n < \bar{Y}_n) \in O(e^{-n})$, i.e. it converges exponentially, while the mean $\bar{X}_n \in O(n^{-0.5})$ converges polynomially only.

   Therefore optimization procedures based on the comparison of configurations (ordinal optimization) rather than the value of the objective function (cardinal optimization) should use samples more efficiently.

2. **Adaptive Allocation**
   [Rinott 1978] developed a two-stage selection procedure that allocates more samples to systems with higher uncertainty instead of sampling all systems equally. This approach can drastically reduce the number of required samples to achieve a given level of confidence about the selection of the best out of several systems. Further improvements were made later on by allocating more samples of a given budget of overall samples to configurations that have a higher probability of being the best.

   The efficiency of samples is increased by allocating them adaptively to more uncertain and more important configurations.

3. **Sequential Analysis**
   [Wald 1947] developed statistical tests that sequentially decide, based on the observations so far, if further samples need to be taken in order to select a hypothesis with a given level of confidence. Therefore the number of samples is not predetermined, but dependent on the observations. The sequential test generally achieves the same accuracy as a test having a fixed number of samples, but using on average only half the number of samples.

   In principal, allowing for a variable number of samples to select the best system increases the efficiency on average.

4. **Bayesian Approach** ([Chick and Inoue 2001b])
   Traditional selection procedures ensure a probability of correct selection within an indifference zone. To achieve this, samples need to be allocated among the alternatives as if the least favorable situation is encountered, i.e. all inferior systems differ exactly by the indifference zone from the best system. In practice a selection procedure seldom encounters this situation. A procedure based on the Bayesian approach would allocate samples for the most probable situation (with respect to some prior distribution) and should therefore be more effective. The drawback is that the procedure cannot guarantee a probability of cor-

rect selection. On the other hand, the Bayesian approach allows to integrate prior information into the procedure, which eventually might improve efficiency.

5. **Using noise**
   Stochastic optimization algorithms deliberately introduce noise into the search process, which allows to escape from local optima. In stochastic environments the objective function is disturbed by noise. So there should be a chance to use the noise from the problem to replace some of the noise introduced by the procedure, therefore reducing the need for noise-reduction by sampling. Additionally, [Fitzpatrick and Grefenstette 1988] observed that moderate noise does not hurt the performance of an EA significantly, making EA robust against noise to a certain degree.

6. **Variance Reduction Techniques**
   Techniques to increase the accuracy of estimation, reduce the estimator's variance by introducing correlation between the samples of a system. This is achieved by using variance reduction techniques like quasi-random numbers (see e.g. [Niederreiter 1992]) or antithetic random variables. The increase in accuracy strongly depends on the structure of the performance criteria.

   In general, when comparing estimations of independent systems' performances, using common random numbers positively correlates the estimations and therefore allows for an increased comparability with the same number of samples.

   Both approaches have the necessity to manipulate the source of randomness in the performance estimation, in general the random number generator.

Except for the last, the above approaches can not be combined with traditional optimizations procedures like Stochastic Programming, because they are based on the objective function rather than comparison of configurations,

they determine the performance on an a-priori fixed number of samples and they are deterministic optimization procedures.

Whereas each of these ideas can be combined with EA: In general EA are rank-based, i.e. they compare configurations only, so they belong to the class of ordinal optimization procedures leading to an improved convergence behavior over cardinal optimization regarding the number of samples. Furthermore this allows for adaptive and sequential selection procedures that can be used within EA.

As EA can not give a performance guarantee for the solution found anyway, the used selection procedures do not need to guarantee a level of confidence, either. Therefore EA are not affected by this Bayesian procedures' disadvantage, but benefit from their increased efficiency. The repeated use of selection procedures within the run of an EA gives the opportunity to acquire information that can be used as prior information for the Bayesian procedures.

EA are stochastic optimization procedures that introduce randomness at several stages. Besides their robustness against moderate noise they give the possibility to use at least some of the randomness originating from the performance estimation of configurations to replace the desired randomness.

All of the ideas – including variance reduction techniques – can be used simultaneously in EA, therefore EA should have the potential to solve real-world stochastic optimization problems.

This thesis develops the tools for applying EA in stochastic environments by deriving procedures that incorporate most of the ideas given above. The achievements do not only advance EA, but also statistical ranking and selection and other ordinal optimization procedures like simulated annealing. The methods presented improve the existing ones by orders of magnitude.

All methods are constructed to deliver a given accuracy with less effort than the existing ones. The question on how much effort – in terms of number of samples – should be spent to increase the accuracy within a single iteration of EA (a generation) versus the number of iterations executed overall, remains still open. Nevertheless most existing choices of accuracy within EA on stochastic optimization problems can be achieved more efficiently with the

methods presented in this thesis.

Improvements not considered here are variance reduction techniques. Using common random numbers for the performance estimation within EA will improve the efficiency further, while quasi-random numbers should not be used as they lead to an underestimation of the variance and therefore deteriorate the methods presented here.

1.4 Structure

This thesis is structured as follows. For ordinal optimization procedures comparison of potential solutions is an elementary operation. Statistical selection procedures provide efficient comparison mechanisms. An overview of existing procedures for statistical selection is given in Chapter 2. The procedures are examined in a comprehensive study and ideas 2-4 from above are integrated, improving the performance of the best known procedures so far.

Chapter 3 shows how one of the most efficient procedures for statistical selection can be integrated into different variants of EA, combining statistical ranking and selection with ordinal optimization. All popular rank-based selection and replacement operators are addressed to efficiently perform optimization on stochastic problems. Additionally, theoretical conditions for the optimal choice of stopping parameters are derived.

In Chapter 4, noise originating from the stochastic problem is used to partly replace the randomization that EA and Simulated Annealing usually introduce into the search process. For that, the applied selection probabilities are deliberately modified, resulting in higher accuracy for a given number of samples or less samples for a given accuracy. The savings in the number of samples are quantified for different parameter settings.

The thesis concludes with a summary of the most important results and points out promising areas of future work.

Chapter 2

Selecting a Selection Procedure

Selection procedures are used in a variety of applications to select the best of a finite set of alternatives. 'Best' is defined with respect to the largest mean, but the mean is inferred with statistical sampling, as in simulation optimization. There is a wide variety of procedures, which gives rise to the question of which selection procedure to select. The main contribution of this chapter is to identify the most effective selection procedures when samples are independent and normally distributed. We also (a) summarize the main structural approaches to deriving selection procedures, (b) derive new procedures, (c) formalize new stopping rules for them, (d) identify strengths and weaknesses of the procedures, and (e) present an innovative empirical test bed with the most extensive numerical comparison of selection procedures to date. The most efficient and easiest to control procedures allocate samples with a Bayesian model for uncertainty about the means, and use new adaptive stopping rules proposed here.

2.1 Overview

Selection procedures are intended to select the best of a finite set of alternatives, where best is determined with respect to the largest sampling mean, but the mean must be inferred via statistical sampling ([Bechhofer, Santner, and Goldsman 1995]). Selection procedures can inform managers how

to select the best of a small set of alternative actions the effects of which are evaluated by simulation ([Nelson and Goldsman 2001]) and are implemented in commercial simulation products like ARENA ([Kelton, Sadowski, and Sadowski 1998]). Selection procedures have also attracted interest in combination with tools like multiple attribute utility theory ([Butler, Morrice, and Mullarkey 2001]), evolutionary algorithms ([Branke and Schmidt 2004]), and discrete optimization via simulation ([Boesel, Nelson, and Kim 2003a]).

Three main approaches to solving the selection problem are distinguished by their assumptions about how the evidence for correct selection is described and sampling allocations are made: the *indifference zone* (IZ, [Kim and Nelson 2005]), the expected *value of information procedure* (VIP, [Chick and Inoue 2001b]), and the *optimal computing budget allocation* (OCBA, [Chen 1996]) approaches. Each approach offers a number of different sampling assumptions, approximations, stopping rules and parameters that combine to define a procedure. With so many variations, the question arises of how to select a selection procedure. The question is important because the demands placed upon simulation optimization algorithms to assist system design choices are increasing.

Only few papers thoroughly assess how those variations compare with each other, but there are a few exceptions. Special cases of the VIP outperform specific IZ and OCBA procedures (in a comparison of two-stage procedures), and specific sequential VIP and OCBA procedures are more efficient than two-stage procedures ([Inoue, Chick, and Chen 1999]). [He, Chick, and Chen 2005] derive an OCBA-type procedure, $\mathcal{OCBA}_{\mathrm{LL}}$, that uses an expected opportunity cost (EOC) loss function, then show that the original $\mathcal{OCBA}$ procedure, the new $\mathcal{OCBA}_{\mathrm{LL}}$ and the VIP-based $\mathcal{LL}$ (described below) perform better than some other procedures in several empirical tests. Procedure $\mathcal{KN}{+}{+}$ is very effective among IZ procedures ([Kim and Nelson 2005]). But even those papers study a limited number of procedures with respect to a limited experimental testbed.

This chapter summarizes the main approaches to selection procedures, derives new procedures and formalizes new stopping rules for the VIP and

OCBA procedures, then addresses the unmet need for an extensive comparison of the new procedures and the top existing IZ, VIP and OCBA procedures. Each procedure makes approximations, and none provides an optimal solution, so it is important to understand the strengths and weaknesses of each approach. Section 2.3 describes new measurements to evaluate each with respect to:

- Efficiency: The mean evidence for correct selection as a function of the mean number of samples.
- Controllability: The ease of setting a procedure's parameters to achieve a targeted evidence level (as opposed to a potentially conservative guarantee that the targeted evidence level is exceeded).
- Robustness: The dependency of a procedure's effectiveness on the underlying problem characteristics.
- Sensitivity: The effect of the parameters on the mean number of samples needed.

The procedures are compared empirically on a large variety of selection problems described in Section 2.4. The test bed is unique not only because of its size, but also its inclusion of randomized problem instances, which may be more realistic in practice than the usual structured problem instances found in the literature.

Each procedure is compared for each metric when applied to broad classes of selection problems described in Section 2.4. The focus here is on applications where the samples are jointly independent and normally distributed with unknown and potentially different variances, or are nearly normally distributed as is the case in stochastic simulation with batching ([Law and Kelton 2000]).

Section 2.5 indicates that Procedure $\mathcal{KN}{+}{+}$ is more efficient than the original VIP and OCBA formulations, but appears to be difficult to control. Certain VIP ($\mathcal{LL}$) and OCBA ($\mathcal{OCBA}$ and $\mathcal{OCBA}_{\mathrm{LL}}$) procedures, when used with new stopping rules below, appear to improve further with respect to

efficiency, controllability and robustness. Section 2.6 addresses further issues that may be important when selecting a selection procedure.

2.2 The Procedures

We first formalize the problem, summarize assumptions and establish notation. Section 2.2.1 describes measures of the evidence of correct selection and, based thereon, introduces new stopping rules that improve efficiency. Sections 2.2.2-2.2.4 describe existing and new variations on sequential procedures from the IZ, VIP and OCBA approaches.

The best of k simulated systems is to be identified, where 'best' means the largest output mean. Analogous results hold if smallest is best. Let X_{ij} be a random variable whose realization x_{ij} is the output of the j-th simulation replication of system i, for $i = 1, \ldots, k$ and $j = 1, 2, \ldots$. Let μ_i and σ_i^2 be the unknown mean and variance of simulated system i, and let $\mu_{[1]} \leq \mu_{[2]} \leq \ldots \leq \mu_{[k]}$ be the ordered means. In practice, the ordering $[\cdot]$ is unknown, and the best system, system $[k]$, is to be identified by simulation. The procedures considered below are derived from the assumption that simulation output is independent and normally distributed, *conditional* on μ_i and σ_i^2, for $i = 1, \ldots, k$.

$$\{X_{ij} : j = 1, 2, \ldots\} \overset{iid}{\sim} \texttt{Normal}\left(\mu_i, \sigma_i^2\right)$$

Although the normality assumption is not always valid, it is often possible to batch a number of outputs so that normality is approximately satisfied. Vectors are written in boldface, such as $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_k)$ and $\boldsymbol{\sigma^2} = (\sigma_1^2, \ldots, \sigma_k^2)$. A problem instance, denoted *configuration* in this chapter, is given by $\boldsymbol{\chi} = (\boldsymbol{\mu}, \boldsymbol{\sigma^2})$.

Let n_i be the number of replications for system i run so far. Let $\bar{x}_i = \sum_{j=1}^{n_i} x_{ij}/n_i$ be the sample mean and $s_i^2 = \sum_{j=1}^{n_i}(x_{ij} - \bar{x}_i)^2/(n_i - 1)$ be the sample variance[1]. Let $\bar{x}_{(1)} \leq \bar{x}_{(2)} \leq \ldots \leq \bar{x}_{(k)}$ be the ordering of the sample means based on all replications seen so far. Equality occurs with probability

[1] In fact we used the numerically more stable variant $s_i^2[n_i] = \frac{1}{n_i - 1}\Sigma_i^2[n_i]$, where $\Sigma_i^2[n_i] = \Sigma_i^2[n_i - 1] + \frac{n_i}{n_i - 1}(x_{in_i} - \bar{x}_i[n_i])^2$ and $\bar{x}_i[n_i], s_i^2[n_i]$ denote mean and variance of the first n_i outputs of system i.

0 in contexts of interest here. (i) denotes the *observed* ordering based on the observed means $\bar{x}_{\cdot}$, while $[i]$ denotes the *true* ordering of systems based on the unknown true means $\mu_{\cdot}$. The quantities n_i, $\bar{x}_i, s_i^2$ and (i) may change as more replications are observed.

Each selection procedure generates estimates $\hat{\mu}_i$ of μ_i, for $i = 1, \ldots, k$. For the procedures studied here, $\hat{\mu}_i = \bar{x}_i$, and a correct selection occurs when the selected system, system $\mathfrak{D}$, is the best system, $[k]$. Usually the system $\mathfrak{D} = (k)$ with the best observed mean $\bar{x}_{\cdot}$ is selected as best, although Procedure $\mathcal{KN}{+}{+}$ below might rarely choose a system that does not have the best sample mean, due to screening. The Student t distribution with mean μ, precision $\kappa = \sigma^{-2}$, and ν degrees of freedom is denoted $\mathtt{St}(\mu, \kappa, \nu)$. If $\nu > 2$ the variance is $\kappa^{-1}\nu/(\nu - 2)$. Denote the cumulative distribution function (cdf) of the standard t ($\mu = 0, \kappa = 1$) and standard Gaussian distribution by $\Phi_\nu(\cdot)$ and $\Phi(\cdot)$ and probability density function (pdf) by $\phi_\nu(\cdot)$ and $\phi(\cdot)$.

2.2.1 Evidence for Correct Selection

This section provides a unified framework for describing both frequentist and Bayesian measures of selection procedure effectiveness and the evidence of correct selection. They are required to derive and compare the procedures below, and are also used within the Bayesian procedures (VIP and OCBA) to decide when the evidence of correct selection is sufficient to stop sampling.

The measures are defined in terms of loss functions. The zero-one loss function, $\mathcal{L}_{0-1}(\mathfrak{D}, \boldsymbol{\mu}) = \mathbb{1}\left\{\mu_{\mathfrak{D}} \neq \mu_{[k]}\right\}$, equals 1 if the best system is not correctly selected, and is 0 otherwise. The zero-one loss function with indifference zone, $\mathcal{L}_{\delta^*}(\mathfrak{D}, \boldsymbol{\mu}) = \mathbb{1}\left\{\mu_{\mathfrak{D}} < \mu_{[k]} - \delta^*\right\}$, relaxes the loss only to 1 if the selected system is not within δ^* of the best. The opportunity cost $\mathcal{L}_{oc}(\mathfrak{D}, \boldsymbol{\mu}) = \mu_{[k]} - \mu_{\mathfrak{D}}$ is 0 if the best system is correctly selected, and is otherwise the difference between the best and selected system. The opportunity cost makes more sense in business applications.

Frequentist Perspective

The IZ procedures take a frequentist perspective. The frequentist probability of correct selection ($\mathrm{PCS}_{\mathrm{iz}}$) is the probability that the system selected as best (system $\mathfrak{D}$) is the system with the highest mean (system $[k]$), conditional on the problem instance. The probability is with respect to the simulation output X_{ij} generated by the procedure (the realizations x_{ij} determine $\mathfrak{D}$).

$$\mathrm{PCS}_{\mathrm{iz}}(\boldsymbol{\chi}) \stackrel{\mathrm{def}}{=} 1 - E\left[\mathcal{L}_{0-1}(\mathfrak{D}, \boldsymbol{\mu}) \,|\, \boldsymbol{\chi}\right] = \Pr\left(\mu_{\mathfrak{D}} = \mu_{[k]} \,|\, \boldsymbol{\chi}\right)$$

Indifference zone procedures attempt to guarantee a lower bound on $\mathrm{PCS}_{\mathrm{iz}}$, subject to the indifference-zone constraint that the best system is at least $\delta^* > 0$ better than the others,

$$\mathrm{PCS}_{\mathrm{iz}}(\boldsymbol{\chi}) \geq 1 - \alpha^*, \text{for all } \boldsymbol{\chi} = (\boldsymbol{\mu}, \boldsymbol{\sigma^2}) \text{ such that } \mu_{[k]} \geq \mu_{[k-1]} + \delta^*. \quad (2.1)$$

[Nelson and Banerjee 2001] showed that some IZ procedures satisfy frequentist probability of good selection guarantees,

$$\mathrm{PGS}_{\mathrm{iz},\delta^*}(\boldsymbol{\chi}) \stackrel{\mathrm{def}}{=} \Pr\left(\mu_{\mathfrak{D}} + \delta^* \geq \mu_{[k]} \,|\, \boldsymbol{\chi}\right) \geq 1 - \alpha^*,$$

for *all* configurations. Let $\mathrm{PICS}_{\mathrm{iz}} = 1 - \mathrm{PCS}_{\mathrm{iz}}$ and $\mathrm{PBS}_{\mathrm{iz},\delta^*} = 1 - \mathrm{PGS}_{\mathrm{iz},\delta^*}$ denote the probability of *incorrect* and *bad* selections.

The frequentist expected opportunity cost ([Chick and Wu 2005]) is

$$\mathrm{EOC}_{\mathrm{iz}}(\boldsymbol{\chi}) \stackrel{\mathrm{def}}{=} E\left[\mathcal{L}_{oc}(\mathfrak{D}, \boldsymbol{\mu}) \,|\, \boldsymbol{\chi}\right] = E\left[\mu_{[k]} - \mu_{\mathfrak{D}} \,|\, \boldsymbol{\chi}\right].$$

Bayesian Perspective

Bayesian procedures use the posterior distribution of the unknown parameters to measure the quality of a selection. Given the data $\mathcal{E}$ seen so far, three

measures of selection quality are

$$\begin{aligned}
\mathrm{PCS}_{\mathrm{Bayes}} &\stackrel{\mathrm{def}}{=} 1 - E\left[\mathcal{L}_{0-1}(\mathfrak{D}, \mathbf{M}) \,|\, \mathcal{E}\right] = \Pr\left(M_{\mathfrak{D}} \geq \max_{i \neq \mathfrak{D}} M_i \,|\, \mathcal{E}\right) \\
\mathrm{PGS}_{\mathrm{Bayes}} &\stackrel{\mathrm{def}}{=} 1 - E\left[\mathcal{L}_{\delta^*}(\mathfrak{D}, \mathbf{M}) \,|\, \mathcal{E}\right] = \Pr\left(M_{\mathfrak{D}} + \delta^* \geq \max_{i \neq \mathfrak{D}} M_i \,|\, \mathcal{E}\right) \\
\mathrm{EOC}_{\mathrm{Bayes}} &\stackrel{\mathrm{def}}{=} E\left[\mathcal{L}_{oc}(\mathfrak{D}, \mathbf{M}) \,|\, \mathcal{E}\right] = E\left[\max_{i=1,\ldots,k} M_i - M_{\mathfrak{D}} \,|\, \mathcal{E}\right], \qquad (2.2)
\end{aligned}$$

the expectation taken over both $\mathfrak{D}$ and the posterior distribution $\mathbf{M}$ of $\boldsymbol{\mu}$ given $\mathcal{E}$. $\mathbf{M}$ is the upper case letter of $\boldsymbol{\mu}$ to indicate the Bayesian perspective that $\boldsymbol{\mu}$ is a realization of its corresponding random variable $\mathbf{M}$. Assuming a noninformative prior distribution for the unknown mean and variance, the posterior marginal distribution for the unknown mean M_i given $n_i > 2$ samples is $\mathtt{St}\left(\bar{x}_i, n_i/s_i^2, \nu_i\right)$ where $\nu_i = n_i - 1$ ([de Groot 1970]). Each Bayesian procedure below selects the system with the best sample mean after all sampling is done, $\mathfrak{D} = (k)$.

Equations 2.2 can be calculated numerically with the following approach: First fix $M_{\mathfrak{D}}$ to a value x, then determine the values for the other M_i and integrate over all x, weighted by the density of $M_{\mathfrak{D}}$. From order statistics we have $\Pr(\max_i X_i \leq x) = \prod_i \Pr(X_i \leq x)$ if the X_i are independent. Let $F_i(\cdot)$ and $f_i(\cdot)$ denote the cdf and pdf of X_i then the cdf and pdf of $\max_i X_i$ are $\prod_i F_i(x)$ and $\sum_i f_i(x) \prod_{j \neq i} F_j(x)$.

$$\begin{aligned}
\mathrm{PGS}_{\mathrm{Bayes}} &= \int_{x=-\infty}^{\infty} \prod_{i \neq \mathfrak{D}} F_i(x + \delta^*) f_{\mathfrak{D}}(x) dx \\
\mathrm{EOC}_{\mathrm{Bayes}} &= \int_{x=-\infty}^{\infty} \int_{y=x}^{\infty} (y - x) \sum_{i \neq \mathfrak{D}} f_i(y) \prod_{j \neq i, \mathfrak{D}} F_j(y) dy f_{\mathfrak{D}}(x) dx \qquad (2.3)
\end{aligned}$$

where $F_i(x) = \Phi_{\nu_i}\left((x - \bar{x}_i)\sqrt{n_i/s_i^2}\right)$, $f_i(x) = \phi_{\nu_i}\left((x - \bar{x}_i)\sqrt{n_i/s_i^2}\right)\sqrt{n_i/s_i^2}$. To have a good selection for a given $M_{\mathfrak{D}} = x$, the other values must be at most $x + \delta^*$, which leads to the expression for $\mathrm{PGS}_{\mathrm{Bayes}}$. Fixing $M_{\mathfrak{D}} = x$ splits $\mathrm{EOC}_{\mathrm{Bayes}}$ in two cases: If $\max_{i \neq \mathfrak{D}} M_i < x$ there is no loss. Otherwise the loss is the expected difference conditional on $\max_{i \neq \mathfrak{D}} M_i \geq x$.

Equations 2.2 can also be calculated with the procedures described in [Genz and Bretz 2002], if the difference of Student variables is assumed to be Student distributed, again. The correlation between the pairwise differences with the selected system $\mathfrak{D}$ is approximately $\frac{s^2_{\mathfrak{D}}/n_{\mathfrak{D}}}{\sqrt{s_i^2/n_i+s^2_{\mathfrak{D}}/n_{\mathfrak{D}}}\sqrt{s_j^2/n_j+s^2_{\mathfrak{D}}/n_{\mathfrak{D}}}}$. The procedures can be even adapted to the case of correlated X_i.

Approximations

The above approaches are computationally very expensive. In the application of selection procedures the effort for high numerical accuracy might be better reduced in return for more time to simulate systems. Therefore we give some approximations needed for the fast calculation of $\mathrm{PGS}_{\mathrm{Bayes}}$ and $\mathrm{EOC}_{\mathrm{Bayes}}$.

Difference of Student variables The Bayesian procedures need to calculate the probability that two independent Student variables $t_i \sim \mathtt{St}\,(\mu_i, \kappa_i, \nu_i)$, $t_j \sim \mathtt{St}\,(\mu_j, \kappa_j, \nu_j)$ differ by at least δ^*. There are several ways to calculate this probability:

- **Numerical Integration**:

$$\begin{aligned}\Pr(t_i+\delta^* < t_j) &= \int_{-\infty}^{\infty} \left(1-\Phi_{\nu_j}\left((x-\mu_j)\sqrt{\kappa_j}\right)\right) \\ &\qquad \phi_{\nu_i}\left((x-\mu_i-\delta^*)\sqrt{\kappa_i}\right)\sqrt{\kappa_i}\,dx \\ &= \int_0^1 \Phi_{\nu_j}\left(-(\Phi^{-1}_{\nu_i}(t)\sqrt{\kappa_i}+\delta^*+\mu_i-\mu_j)\sqrt{\kappa_j}\right)\,dt\end{aligned} \tag{2.4}$$

- **Welch's approximation** ([Welch 1938]) for the degrees of freedom of the difference (see [Law and Kelton 2000], p. 599):
  $t_i - t_j \approx \mathtt{St}\left(\mu_i-\mu_j, (\kappa_i^{-1}+\kappa_j^{-1})^{-1}, \nu_{ij}\right)$, where
  $\nu_{ij} = (\kappa_i^{-1}+\kappa_j^{-1})^2/(\kappa_i^{-2}/\nu_i + \kappa_j^{-2}/\nu_j)$

$$\Pr(t_i+\delta^* < t_j) \approx \Phi_{\nu_{ij}}(-d^*_{ij}) \tag{2.5}$$

  with $d^*_{ij} = (\delta^*+\mu_i-\mu_j)\sqrt{\kappa_i^{-1}+\kappa_j^{-1}}$ set for convenience. Our experi-

ments indicate that this approximation is a lower bound, although we could not prove this observation.

- **Wilson's approximation** ([Wilson and Pritsker 1984]) instead of Welch's for the degrees of freedom as above:
  $\nu'_{ij} = (\kappa_i^{-1} + \kappa_j^{-1})^2 / (\kappa_i^{-2}/(\nu_i + 2) + \kappa_j^{-2}/(\nu_j + 2)) - 2$

$$\Pr(t_i + \delta^* < t_j) \approx \Phi_{\nu'_{ij}}(-d^*_{ij}) \tag{2.6}$$

- **Gaussian approximation** of the difference:
  $t_i - t_j \approx \mathcal{N}(\mu_i - \mu_j, \kappa_i^{-1} + \kappa_j^{-1})$

$$\Pr(t_i + \delta^* < t_j) \approx \Phi(-d^*_{ij}) \tag{2.7}$$

- **Chernoff's inequality** for the Gaussian-approximation:
  $\Phi(x) \leq \exp(-x^2/2)$ for $x < 0$

$$\Pr(t_i + \delta^* < t_j) \approx \Phi\left(-d^*_{ij}\right) = \begin{cases} \mu_i > \mu_j : & \leq \exp(-{d^*_{ij}}^2/2) \\ \mu_i < \mu_j : & \geq 1 - \exp(-{d^*_{ij}}^2/2) \end{cases} \tag{2.8}$$

Bonferroni's Bound Bonferroni's inequality (see [Law and Kelton 1991]) gives $\Pr(\bigcap_{i=1}^k X_i < 0) \geq 1 - \sum_{i=1}^k [1 - \Pr(X_i < 0)]$, which can be applied to give a lower bound for $\mathrm{PCS}_{\mathrm{Bayes}}$

$$\mathrm{PCS}_{\mathrm{Bayes}} \geq 1 - \sum_{j:(j)\neq(k)} [1 - \Pr\left(M_{(k)} < M_{(j)} \,|\, \mathcal{E}\right)].$$

Approximating the right hand side with Welch's approximation defines

$$\mathrm{PCS}_{\mathrm{Bonf}} = 1 - \sum_{j:(j)\neq(k)} \Phi_{\nu_{(j)(k)}}(d^*_{jk}),$$

which might deliver negative values for $\mathrm{PCS}_{\mathrm{Bonf}}$.

The term $\mathrm{EOC}_{\mathrm{Bayes}}$ may be expensive to compute if $k > 2$. Summing

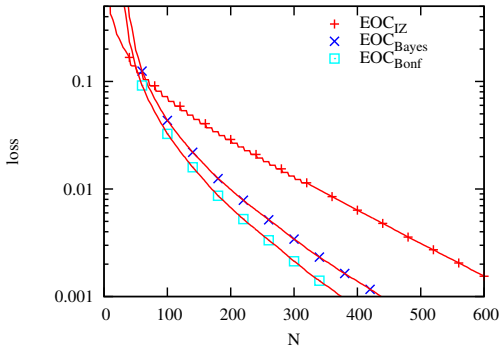

Figure 2.1: Different measures and approximations for correct selection (Equal samples for $k = 10$ systems with a pairwise difference of $\delta = 0.5$ and equal variances).

the losses from $(k-1)$ pairwise comparisons between the current best and each other system gives an easily computed upper bound ([Chick and Inoue 2001b; Chick and Inoue 2002]). Let $f_{(j)(k)}(\cdot)$ be the posterior pdf for the difference $M_{(j)} - M_{(k)}$ given all data $\mathcal{E}$ (approximately $\mathtt{St}\left(-d_{(j)(k)}, \lambda_{jk}, \nu_{(j)(k)}\right)$ distributed), and set

$$\Psi_\nu[s] = \int_{u=s}^{\infty} (u-s)\phi_\nu(u)du = \frac{\nu+s^2}{\nu-1}\phi_\nu(s) - s\Phi_\nu(-s). \tag{2.9}$$

Then

$$\begin{aligned} \mathrm{EOC}_{\mathrm{Bayes}} &\leq \sum_{j:(j)\neq(k)} \int_{w=0}^{\infty} w\, f_{(j)(k)}(w)\, dw \\ &\approx \sum_{j:(j)\neq(k)} \lambda_{jk}^{-1/2} \Psi_{\nu_{(j)(k)}}\left[d_{jk}^*\right] = \mathrm{EOC}_{\mathrm{Bonf}}. \end{aligned} \tag{2.10}$$

The deviation between the Frequentist loss ($\mathrm{EOC}_{\mathrm{iz}}$), the Bayesian loss ($\mathrm{EOC}_{\mathrm{Bayes}}$) and its approximation with the Bonferroni-like bound is shown in Figure 2.1. The two definitions of loss differ clearly. Although Bonferroni's bound is an upper bound for the Bayesian loss, $\mathrm{EOC}_{\mathrm{Bonf}}$ is below $\mathrm{EOC}_{\mathrm{Bayes}}$, which therefore must be originated in Welch's approximation.

Slepian's Bound Approximations in the form of bounds on the above losses are useful to derive sampling allocations and to define stopping rules. *Slepian's* inequality ([Tong 1980]) states the posterior evidence that system (k) is best satisfies

$$\mathrm{PCS}_{\mathrm{Bayes}} \geq \prod_{j:(j)\neq(k)} \Pr\left(M_{(k)} \geq M_{(j)} \,|\, \mathcal{E}\right). \tag{2.11}$$

The right hand side of Inequality (2.11) is approximately (Welch)

$$\mathrm{PCS}_{\mathrm{Slep}} = \prod_{j:(j)\neq(k)} \Phi_{\nu_{(j)(k)}}(d^*_{jk}),$$

where d^*_{jk} is the normalized distance for systems (j) and (k),

$$\begin{aligned} d^*_{jk} &= d_{(j)(k)}\lambda_{jk}^{1/2} \text{ with } d_{(j)(k)} = \bar{x}_{(k)} - \bar{x}_{(j)} \text{ and } \lambda_{jk}^{-1} = \frac{s^2_{(j)}}{n_{(j)}} + \frac{s^2_{(k)}}{n_{(k)}}, \\ \nu_{(j)(k)} &= \frac{[s^2_{(j)}/n_{(j)} + s^2_{(k)}/n_{(k)}]^2}{[s^2_{(j)}/n_{(j)}]^2/(n_{(j)}-1) + [s^2_{(k)}/n_{(k)}]^2/(n_{(k)}-1)}. \end{aligned}$$

$\mathrm{PCS}_{\mathrm{Slep}} \geq \mathrm{PCS}_{\mathrm{Bonf}}$, where equality only holds for $k = 2$ and $\mathrm{PCS}_{\mathrm{Slep}}$ is strictly tighter than $\mathrm{PCS}_{\mathrm{Bonf}}$ for more than 2 systems.[2] We therefore mostly used $\mathrm{PCS}_{\mathrm{Slep}}$.

The following variation incorporates an indifference zone parameter δ^* to approximate the probability that the difference between the selected system and the true best system is no more than δ^* (PGS for probability of good selection). Note that $\mathrm{PCS}_{\mathrm{Slep}} = \mathrm{PGS}_{\mathrm{Slep},0}$.

$$\mathrm{PGS}_{\mathrm{Slep},\delta^*} = \prod_{j:(j)\neq(k)} \Phi_{\nu_{(j)(k)}}(\lambda_{jk}^{1/2}(\delta^* + d_{(j)(k)})). \tag{2.12}$$

[2] For $k = 2$, $\mathrm{PCS}_{\mathrm{Bonf}}$ and $\mathrm{PCS}_{\mathrm{Slep}}$ are equal. For 3 systems, if the additional $\Phi_\nu(\cdot) = 1$ then $\mathrm{PCS}_{\mathrm{Bonf}} = \mathrm{PCS}_{\mathrm{Slep}}$. As $\Phi_\nu(\cdot)$ is always below 1 and $\mathrm{PCS}_{\mathrm{Bonf}}$ decreases faster than $\mathrm{PCS}_{\mathrm{Slep}}$ for lower $\Phi_\nu(\cdot)$, $\mathrm{PCS}_{\mathrm{Bonf}} < \mathrm{PCS}_{\mathrm{Slep}}$ holds for all $k > 2$.

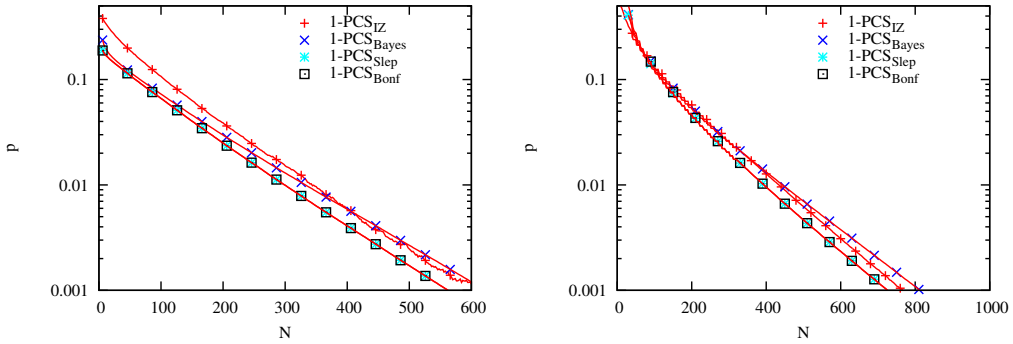

Figure 2.2: Different measures and approximations for correct selection (Equal samples for $k = 2$ systems with a difference of $\delta = 0.25$ (left) and $k = 10$ with a difference $\delta = 0.5$ of neighboring systems (right). The variances of all systems are equal).

[Chen and Kelton 2005] used max instead of +,

$$\mathrm{PCS}_{\mathrm{Slep},\delta^*} = \prod_{j:(j)\neq(k)} \Phi_{\nu_{(j)(k)}}(\lambda_{jk}^{1/2} \max\{\delta^*, d_{(j)(k)}\}),$$

which can be interpreted as a maximum-likelihood-like estimator for $\mathrm{PGS}_{\mathrm{iz},\delta^*}$ conditional on $\mu_{(k)} \geq \mu_{(j)} + \delta^*$ for all $(j) \neq (k)$.

The deviation between $\mathrm{PCS}_{\mathrm{iz}}$, $\mathrm{PCS}_{\mathrm{Bayes}}$, $\mathrm{PCS}_{\mathrm{Slep}}$ and $\mathrm{PCS}_{\mathrm{Bonf}}$ is shown in Figure 2.2. Even for 10 systems, there is no visible difference between $\mathrm{PCS}_{\mathrm{Slep}}$ and $\mathrm{PCS}_{\mathrm{Bonf}}$. The difference is only significant for a few samples. Obviously $\mathrm{PCS}_{\mathrm{Slep}}$ is not a lower bound for $\mathrm{PCS}_{\mathrm{Bayes}}$, which further supports the claim, that Equation 2.5 underestimates $\Pr(t_i + \delta^* < t_j)$.

Stopping Rules

The VIP and OCBA procedures defined below will make use of $\mathrm{EOC}_{\mathrm{Bonf}}$ and $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$ to decide when to stop sampling. In particular, the following **stopping rules** are considered:

1. Sequential ($\mathcal{S}$): Repeat sampling while $\sum_{i=1}^{k} n_i < B$ for some specified total budget B.

2. Probability of good selection ($\mathrm{PGS}_{\mathrm{Slep},\delta^*}$): Repeat while $\mathrm{PGS}_{\mathrm{Slep},\delta^*} <$

$1 - \alpha^*$ for a specified probability target $1 - \alpha^*$ and given $\delta^* \geq 0$.

3. Expected opportunity cost (EOC_{Bonf}): Repeat while $\text{EOC}_{\text{Bonf}} > \beta^*$, for a specified EOC target β^*.

The IZ requires $\delta^* > 0$, but we allow $\delta^* = 0$ for the VIP and OCBA to allow for a pure PCS-based stopping condition. We use PCS_{Slep} to denote $\text{PGS}_{\text{Slep},0}$. Previously published sequential VIP and OCBA work used the $\mathcal{S}$ stopping rule. The other stopping rules will be shown to improve the efficiency of both approaches.

2.2.2 Indifference Zone (IZ)

The IZ approach ([Bechhofer, Santner, and Goldsman 1995] seeks to guarantee $\text{PCS}_{\text{iz}} \geq 1 - \alpha^*$, whenever the best system is at least δ^* better than the other systems. The indifference-zone parameter δ^* is typically elicited as the smallest difference in mean performance that is significant to the decision-maker.

Early IZ procedures were statistically conservative in the sense of excess sampling unless unfavorable configurations of the means were found. The $\mathcal{KN}$ family of procedures improves sampling efficiency over a broad set of configurations [Kim and Nelson 2001]. While a PCS guarantee in the sense of Equation (2.1) was not proven, an asymptotic guarantee as $\delta^* \to 0$ was shown. One member of the family, $\mathcal{KN}$++ [Goldsman, Kim, Marshall, and Nelson 2002], might be considered to be the state of the art for the IZ approach.

Kim Nelson That procedure can handle the more general case of correlated simulation output. Here we specialize Procedure $\mathcal{KN}$++ for *independent* output. The procedure screens out some systems as more runs are observed, and each noneliminated system is simulated the same number of times. We used $\xi = 1$ replication per stage per noneliminated system and sample variance updates in every stage.

Procedure $\mathcal{KN}$++ (independent samples)

1. Specify a confidence level $1-\alpha^* > 1-1/k$, an indifference-zone parameter $\delta^* > 0$, a first-stage sample size $n_0 > 2$ per system, and a number ξ of samples to run per noneliminated system per subsequent stage.

2. Initialize the set of noneliminated systems, $I \leftarrow \{1, \ldots, k\}$, set $n \leftarrow 0, \tau \leftarrow n_0$, $\beta \leftarrow 1 - (1-\alpha^*)^{1/(k-1)}$.

3. WHILE $|\, I \,| > 1$ DO another stage:

   (a) Observe τ additional samples from system i, for all $i \in I$. Set $n \leftarrow n + \tau$. Set $\tau \leftarrow \xi$.

   (b) Update: Set $\eta \leftarrow \frac{1}{2}\left[(2\beta)^{-2/(n-1)} - 1\right]$ and $h^2 \leftarrow 2\eta(n-1)$. For all $i \in I$, update the sample statistics $\bar{x}_i$ and s_i^2.

   (c) Screen: For all $i, j \in I$ and $i > j$, set $d_{ij} \leftarrow \bar{x}_j - \bar{x}_i$ and $\epsilon_{ij} \leftarrow \max\left\{0, \frac{\delta^*}{2n}\left(\frac{h^2(s_i^2+s_j^2)}{\delta^{*2}} - n\right)\right\}$. If $d_{ij} > \epsilon_{ij}$ then $I \leftarrow I \backslash \{i\}$. If $d_{ij} < -\epsilon_{ij}$ then $I \leftarrow I \backslash \{j\}$.

4. Return remaining system, system $\mathfrak{D}$, as best.

For correlated samples, e.g. when using common random numbers between systems, $s_i^2 + s_j^2$ is replaced by the estimated common variance of the difference $\frac{1}{n-1}\sum_{k=1}^{n}(x_{ik} - x_{jk} - \bar{x}_i + \bar{x}_j)^2$.

In the literature some variations on $\mathcal{KN}$++ are found. In [Kim and Nelson 2001] $\beta = \alpha^*/(k-1)$ is used. [Goldsman, Kim, Marshall, and Nelson 2002] replaced η by $-\ln(2\beta)/(n-1)$, which is approximately equal to the η chosen above.

Wald's Sequential Probability Ratio Test [Wald 1947] was the first to use sequential sampling in connection with Hypotheses tests. Wald's Sequential Probability Ratio Test (SPRT) is designed to decide between simple hypotheses. The test stops, if the ratio of the probabilities under the hypotheses exceeds given bounds. For a random variable X depending on a parameter θ let the null hypothesis $H_0 : \theta = \theta_0$, the alternative hypothesis $H_1 : \theta = \theta_1$, the error probability of type I (wrongly rejecting H_0) equal α and the error probability of type II (wrongly rejecting H_1) β.

The test decides, which hypothesis to select on the basis of observations $x_1, x_2, \ldots$. More specifically, it continues sampling as long as

$$B < \frac{f_1(x_1, x_2, \ldots)}{f_0(x_1, x_2, \ldots)} < A \tag{2.13}$$

for given constants A, B and $f_i(\cdot)$ being the pdf for all observations under H_i. If the ratio is $\leq B$ then H_0 is selected, if $\geq A$, H_1 is selected. Wald shows that this test minimizes the mean number of samples for both $\theta = \theta_0$ and $\theta = \theta_1$ approximately and guarantees that the error probabilities are below α and β, if $A = \frac{1-\beta}{\alpha}$ and $B = \frac{\beta}{1-\alpha}$.

[Baum and Veeravalli 1994] extend the SPRT for multiple disjoint hypothesis. The MSPRT selects hypothesis H_i, if

$$1 - \frac{f_i(x_1, x_2, \ldots)}{\sum_{j=1..k} f_j(x_1, x_2, \ldots)} \quad \leq \quad \alpha_i, \tag{2.14}$$

where α_i is the error probability for erroneously selecting H_i. MSPRT equals Wald's SPRT for $k = 2$ systems.

To apply the SPRT for selection, we can test the hypothesis that system i is at least δ^* better than all other systems, i.e. $H_i : \forall j \neq i : \mu_i \geq \mu_j + \delta^*$. The difference of the observed means under H_i is approximately $\mathtt{St}\left(\delta_{ij}, (s_i^2/n_i + s_j^2/n_j)^{-1}, \nu_{ij}\right)$-distributed, where ν_{ij} is approximated with Welch and $\delta_{ij} \geq \delta^*$. Assuming independence of the observed means, $f_i(x_1, x_2, \ldots)$ can be approximated by

$$f_i(x_1, x_2, \ldots) \quad \approx \quad \prod_{j \neq i} \phi_{\nu_{ij}} \left(\frac{\bar{x}_i - \bar{x}_j - \delta_{ij}}{\sqrt{s_i^2/n_i + s_j^2/n_j}} \right) \frac{1}{\sqrt{s_i^2/n_i + s_j^2/n_j}}. \tag{2.15}$$

We set $\delta_{ij} = \max\{\delta^*, \bar{x}_i - \bar{x}_j\}$ in the calculation of the selection criterion, so the pdf used is the maximal pdf conditional on the hypothesis H_i:

$\max_{\delta_{ij} \geq \delta^*} f_i(\cdot)$. We further set $\alpha_i = \alpha$ equally. The stopping rule is then

$$1 - \frac{\prod_{j \neq (k)} \phi_{\nu_{(k)j}} \left(\frac{\max\{\bar{x}_{(k)} - \bar{x}_j - \delta^*, 0\}}{\sqrt{s^2_{(k)}/n_{(k)} + s^2_j/n_j}} \right)}{\sum_{i=1..k} \prod_{j \neq i} \phi_{\nu_{ij}} \left(\frac{\max\{\bar{x}_i - \bar{x}_j - \delta^*, 0\}}{\sqrt{s^2_i/n_i + s^2_j/n_j}} \right)} \quad \leq \quad \alpha \tag{2.16}$$

and the system with the best mean $\bar{x}_k$ is the selected. MSPRT can be used as a stopping rule for any of the Bayesian procedures or as a variant for allocation in an OCBA-like procedure.

2.2.3 Value of Information Procedure (VIP)

Two VIPs in [Chick and Inoue 2001b] allocate samples to each alternative in order to maximize the expected value of information (EVI) of the samples, subject to a sampling budget constraint. Procedures 0-1($\mathcal{S}$) and $\mathcal{LL}(\mathcal{S})$ are sequential variations of those procedures that improve Bonferroni bounds for $\text{PCS}_{\text{Bayes}}$ (the expected 0-1 loss) and $\text{EOC}_{\text{Bayes}}$ ($\mathcal{LL}$ for linear loss), respectively. They allocate τ replications per stage until a total of B replications are run.

While the original stopping rule allows for full control of the number of replications, the procedures outlined below allow to choose from any of the stopping rules defined in Section 2.2.1. As we will demonstrate later, alternative stopping rules make the procedure significantly more efficient. Furthermore, new procedures are introduced that use different approximations for the EVI of samples.

Procedure 0-1.

1. Specify a first-stage sample size $n_0 > 2$, and a total number of samples $\tau > 0$ to allocate per subsequent stage. Specify stopping rule parameters.

2. Run independent replications $X_{i1}, \ldots, X_{in_0}$, and initialize the number of replications $n_i \leftarrow n_0$ run so far for each system, $i = 1, \ldots, k$.

3. Determine the sample statistics $\bar{x}_i$ and s_i^2, and the order statistics, so that $\bar{x}_{(1)} \leq \ldots \leq \bar{x}_{(k)}$.

4. WHILE stopping rule not satisfied DO another stage:

   (a) Initialize set of systems considered for additional replications, $\mathcal{S} \leftarrow \{1, \ldots, k\}$.

   (b) For each (i) in $\mathcal{S} \backslash \{(k)\}$: If $(k) \in \mathcal{S}$ then set $\lambda_{ik}^{-1} \leftarrow s_{(i)}^2/n_{(i)} + s_{(k)}^2/n_{(k)}$, and set $\nu_{(i)(k)}$ with Welch's approximation. If $(k) \notin \mathcal{S}$ then set $\lambda_{ik} \leftarrow n_{(i)}/s_{(i)}^2$ and $\nu_{(i)(k)} \leftarrow n_{(i)} - 1$.

   (c) Tentatively allocate a total of τ replications to systems $(i) \in \mathcal{S}$ (set $\tau_{(j)} \leftarrow 0$ for $(j) \notin \mathcal{S}$):

$$\tau_{(i)} \leftarrow \frac{(\tau + \sum_{j \in \mathcal{S}} n_j)(s_{(i)}^2 \gamma_{(i)})^{\frac{1}{2}}}{\sum_{j \in \mathcal{S}} (s_j^2 \gamma_j)^{\frac{1}{2}}} - n_{(i)},$$

   where

$$\gamma_{(i)} \leftarrow \begin{cases} \lambda_{ik} d_{ik}^* \phi_{\nu_{(i)(k)}}(d_{ik}^*) & \text{for } (i) \neq (k) \\ \sum_{(j) \in \mathcal{S} \backslash \{(k)\}} \gamma_{(j)} & \text{for } (i) = (k). \end{cases}$$

   (d) If any $\tau_i < 0$ then fix the nonnegativity constraint violation: remove (i) from $\mathcal{S}$ for each (i) such that $\tau_{(i)} \leq 0$, and go to Step 4b. Otherwise, round the τ_i so that $\sum_{i=1}^k \tau_i = \tau$ and go to Step 4e.

   (e) Run τ_i additional replications for system i, for $i = 1, \ldots, k$. Update sample statistics $n_i \leftarrow n_i + \tau_i$; $\bar{x}_i$; s_i^2, and the order statistics, so $\bar{x}_{(1)} \leq \ldots \leq \bar{x}_{(k)}$.

5. Select the system with the best estimated mean, $\mathfrak{D} = (k)$.

The formulas in Step 4b use the Welch approximation, and the formulas in Step 4c are derived from optimality conditions to improve a Bonferroni-like bound on the EVI for asymptotically large τ ([Chick and Inoue 2001b]).

Depending on the stopping rule used, the resulting procedures are named 0-1($\mathcal{S}$), 0-1($\text{PGS}_{\text{Slep},\delta^*}$), 0-1($\text{EOC}_{\text{Bonf}}$), with the stopping rule in parentheses. The inclusion of a parameter $\delta^* \geq 0$ for $\text{PGS}_{\text{Slep},\delta^*}$ permits early stopping

if the difference between strongly competing alternatives is negligible. The IZ requires $\delta^* > 0$, but we allow $\delta^* = 0$ for the VIP and OCBA to assess a pure PCS-based stopping condition. If $\delta^* = 0$ we denote the stopping rule PCS_{Slep}, which is algebraically equivalent when $\delta^* = 0$, to emphasize the relation to PCS.

Procedure $\mathcal{LL}$ is a variant of 0-1 where sampling allocations seek to minimize EOC_{Bonf}. The initials $\mathcal{LL}$ (linear loss) are used rather than $\mathcal{OC}$ (opportunity cost) to avoid naming conflicts with $\mathcal{OCBA}$.

Procedure $\mathcal{LL}$. Same as Procedure 0-1, except set $\gamma_{(i)}$ in Step 4c to

$$\gamma_{(i)} \leftarrow \begin{cases} \lambda_{ik}^{1/2} \frac{\nu_{(i)(k)}+(d_{ik}^*)^2}{\nu_{(i)(k)}-1} \phi_{\nu_{(i)(k)}}(d_{ik}^*) & \text{for } (i) \neq (k) \\ \sum_{(j)\in\mathcal{S}\setminus\{(k)\}} \gamma_{(j)} & \text{for } (i) = (k) \end{cases} \tag{2.17}$$

New Small-Sample Procedures. Procedures 0-1 and $\mathcal{LL}$ allocate additional replications using an EVI approximation based on asymptotically *large* number of replications (τ) per stage. That EVI approximation also uses a Bonferroni-like bound, and necessitates the Welch approximation.

An improvement might be obtained by better approximating EVI when there are a *small* number of replications per stage that are all run for one system. The Procedure $\mathcal{LL}_1$ is derived by removing the asymptotic approximation from the derivation of $\mathcal{LL}$, as well as the Bonferroni and Welch approximations. The procedure's name is distinguished from its large-sample counterpart by the subscript $_1$ (1 system gets all replications per stage). The procedure uses the following variables.

$$\begin{aligned} d_{\{jk\}}^* &= \lambda_{\{jk\}}^{1/2} d_{(j)(k)} \\ \lambda_{\{jk\}}^{-1} &= \left(\frac{\tau_{(k)} s_{(k)}^2}{n_{(k)}(n_{(k)}+\tau_{(k)})} + \frac{\tau_{(j)} s_{(j)}^2}{n_{(j)}(n_{(j)}+\tau_{(j)})} \right) \end{aligned} \tag{2.18}$$

Procedure $\mathcal{LL}_1$. Same as Procedure 0-1, except replace Steps 4a-4d by:

- For each $i \in \{1, 2, \ldots, k\}$, see if allocating to (i) is best:
  - Tentatively set $\tau_{(i)} \leftarrow \tau$ and $\tau_\ell \leftarrow 0$ for all $\ell \neq (i)$; set $\lambda_{\{jk\}}^{-1}$, $d_{\{jk\}}^*$

with Equation (2.18) for all j.

– Compute the EVI,

$$\mathrm{EVI}_{LL,(i)} = \begin{cases} \lambda_{\{ik\}}^{-1/2} \Psi_{n_{(i)}-1}\left[d^*_{\{ik\}}\right] & \text{if } (i) \neq (k) \\ \lambda_{\{k-1,k\}}^{-1/2} \Psi_{n_{(k)}-1}\left[d^*_{\{k-1,k\}}\right] & \text{if } (i) = (k). \end{cases}$$

- Set $\tau_{(i)} \leftarrow \tau$ for the system that maximizes $\mathrm{EVI}_{LL,(i)}$, and $\tau_\ell \leftarrow 0$ for the others.

The sampling allocation in Procedure 0-1 is based on two asymptotic approximations. Procedure 0-1_1 avoids one of them, as well as the Bonferroni and Welch approximations.

Procedure 0-1_1. Same as Procedure $\mathcal{LL}_1$, except the EVI is approximated with respect to the 0-1 loss,

$$\mathrm{EVI}_{0-1,(i)} = \begin{cases} \Phi_{n_{(i)}-1}(-d^*_{\{ik\}}) & \text{if } (i) \neq (k) \\ \Phi_{n_{(k)}-1}(-d^*_{\{k-1,k\}}) & \text{if } (i) = (k) \end{cases} \tag{2.19}$$

Variations for the stopping rules are named analogously as summarized in Section 2.2.7.

2.2.4 OCBA Procedures

The OCBA is a class of procedures that was initially proposed by [Chen 1996] and that has several variations (e.g. [Inoue, Chick, and Chen 1999; Chen, Yücesan, Dai, and Chen 2005]). The variations involve different approximations for $\mathrm{PCS}_{\mathrm{Bayes}}$, and different thought experiments for how additional samples might improve the probability of correct selection. Here we specify the idea behind the OCBA and the variations used for this thesis. The OCBA assumes that if an additional τ replications are allocated for system i, but none are allocated for the other systems, then the standard error is scaled back accordingly. The usual OCBA assumes normal distributions to

approximate the effect, but we use Student distributions,

$$\begin{aligned} \tilde{M}_i &\sim \mathtt{St}\left(\bar{x}_i, (n_i+\tau)/s_i^2, n_i - 1 + \tau\right) \\ \tilde{M}_j &\sim \mathtt{St}\left(\bar{x}_j, n_j/s_j^2, n_j - 1\right) \qquad\qquad \text{for } j \neq i, \end{aligned}$$

for consistency with a Bayesian assumption for the unknown σ_i^2. [Chen, Yücesan, Dai, and Chen 2005] and [Branke, Chick, and Schmidt 2005b] found no notable difference in performance when comparing normal versus Student distributions for the $\tilde{M}_i$.

The effect of allocating an additional τ replications to system i, but no replications to the others, leads to an *estimated approximate probability of correct selection* (EAPCS) evaluated with respect to $\tilde{\mathbf{M}} = (\tilde{M}_1, \ldots, \tilde{M}_k)$, and with $\tilde{M}_{(j)} - \tilde{M}_{(k)}$ approximated using Welch's approximation.

$$\begin{aligned} \text{EAPCS}_i &= \prod_{j:(j)\neq(k)} \Pr\left(\tilde{M}_{(j)} < \tilde{M}_{(k)} \,|\, \mathcal{E}\right) \\ &\approx \prod_{j:(j)\neq(k)} (1 - \Phi_{\tilde{\nu}_{(j)(k)}}(\tilde{\lambda}_{jk}^{1/2} d_{(j)(k)})) \\ \tilde{\lambda}_{jk}^{-1} &= \frac{s_{(k)}^2}{n_{(k)} + \tau \mathbb{1}\{(k) = i\}} + \frac{s_{(j)}^2}{n_{(j)} + \tau \mathbb{1}\{(j) = i\}} \end{aligned} \tag{2.20}$$

where $\mathbb{1}\{\cdot\}$ is 1 if the argument is true, and 0 otherwise.

These approximations result in a sequential OCBA algorithm that greedily allocates samples to systems that most increase $\text{EAPCS}_i - \text{PCS}_{\text{Slep}}$ at each stage. An innovation for OCBA here is that sampling continues until a stopping rule from Section 2.2.1 is satisfied.

Procedure $\mathcal{OCBA}$.

1. Specify a first-stage sample size $n_0 > 2$, a number q of systems to simulate per stage, a sampling increment $\tau > 0$ to allocate per subsequent stage, and stopping rule parameters.

2. Run independent replications $X_{i1}, \ldots, X_{in_0}$, and initialize the number of replications $n_i \leftarrow n_0$ run so far for each system, $i = 1, \ldots, k$.

3. Determine the sample statistics $\bar{x}_i$ and s_i^2 and the sample mean ordering, so that $\bar{x}_{(1)} \leq \ldots \leq \bar{x}_{(k)}$.

4. WHILE stopping rule not satisfied DO another stage:

   (a) Compute EAPCS_i for $i = 1, \ldots, k$.

   (b) Set $\tau_i \leftarrow \tau/q$ for the q systems with largest $\text{EAPCS}_i - \text{PCS}_{\text{Slep}}$, set $\tau_j \leftarrow 0$ for the others.

   (c) Run τ_i additional observations from system i.

   (d) For all i with $\tau_i > 0$, update $n_i \leftarrow n_i + \tau_i$, the sample statistics $\bar{x}_i$, s_i^2, and order statistics, so that $\bar{x}_{(1)} \leq \ldots \leq \bar{x}_{(k)}$.

5. Select the system with the best estimated mean, $\mathfrak{D} = (k)$.

[He, Chick, and Chen 2005] proposed an OCBA variation that accounts for the expected opportunity cost. Define AEOC to be the approximation to EOC_{Bonf} in the right hand side of Equation (2.10), and set

$$\text{EEOCS}_i = \sum_{j:(j)\neq(k)} \tilde{\lambda}_{jk}^{-1/2} \Psi_{\tilde{\nu}_{(j)(k)}} \left[\tilde{\lambda}_{jk}^{1/2} d_{(j)(k)} \right].$$

Procedure $\mathcal{OCBA}_{\mathbf{LL}}$ is a variation of $\mathcal{OCBA}$ that allocates replications to systems that maximize the improvement in expected opportunity cost (linear loss), $\text{AEOC} - \text{EEOCS}_i$ in Step 4b.

Yet another OCBA heuristic incorporates the indifference zone parameter δ^* into the sampling allocation (not just the stopping rule). Let $\text{EAPGS}_{i,\delta^*}$ generalize $\text{PGS}_{\text{Slep},\delta^*}$ by computing it with respect to $\tilde{\mathbf{M}}$. Procedure $\mathcal{OCBA}_{\delta^*}$ allocates replications to systems that most improve an estimated probability of a good selection, $\text{EAPGS}_{i,\delta^*} - \text{PGS}_{\text{Slep},\delta^*}$, in Step 4b. This differs from how δ^* was incorporated into OCBA by [Chen and Kelton 2005]. We denote the latter by $\mathcal{OCBA}_{max,\delta^*}$.

All OCBA variations above were implemented as fully sequential procedures here ($q = 1$ and $\tau = 1$).

2.2.5 Optimal allocation

For $k = 2$ systems and equal variances it is optimal to allocate samples equally among the systems. We will now derive an "optimal" allocation scheme. For the derivation, we assume the means and variances to be known and allocate new samples, so that the probability for the best system to have the largest observed mean is maximized. The allocation is determined independently of the observations made, so it's statically. The purpose of this allocation is to give a theoretical bound on the maximal effect, allocation can add to selection procedures.

$$\begin{aligned} \max_{n_1,\ldots,n_k} \quad & \Pr(\bar{X}_{[k]} > \bar{X}_i \, \forall i \neq [k]) \\ s.t. \quad & \sum_{i=1}^{k} n_i = n \\ & n_i \geq 0, \end{aligned} \tag{2.21}$$

where $\bar{X}_i$ is the distribution of the observed means after n_i samples. The probability can be calculated by $\int_{-\infty}^{\infty} \prod_{i \neq k} F_i(x, n_{[i]}) f_k(x, n_{[k]}) dx$, where $F_i(x,n) = \Phi(\frac{x-\mu_{[i]}}{\sigma_{[i]}}\sqrt{n})$ and $f_i(x,n) = \phi(\frac{x-\mu_{[i]}}{\sigma_{[i]}}\sqrt{n})\frac{\sqrt{n}}{\sigma_{[i]}}$ are the cdf and pdf of $\bar{X}_i$ for a given number of samples n.

We implemented optimal allocation by greedily allocating samples for the system that increases the objective at most.

2.2.6 Computational Issues

The implementation that generated the analysis and graphs in this chapter used the Gnu Scientific Libary (gsl) for calculating cdfs, the Mersenne twister random number generator [Matsumoto and Nishimura 1998, (with 2002 revised seeding)] and FILIB++ ([Lerch, Tischler, von Gudenberg, Hofschuster, and Kraemer 2001]) for interval arithmetic. Calculations were run on a mixed cluster of up to 120 nodes. The nodes were running Linux 2.4 and Windows XP with Intel P4 and AMD Athlon processors ranging from 2 to 3 GHz. The program is written in C++ and jobs were distributed with

the JOSCHKA-System ([Bonn, Toussaint, and Schmeck 2005]).

Numerical stability problems may arise in implementations of the OCBA, 0-1_1 and $\mathcal{LL}_1$ allocations even with double-precision floating point arithmetic as the total number of replications gets quite large. For example, $\Psi_\nu[s]$ in Equation (2.9) is strictly positive for finite s but may evaluate to nonpositive values (e.g. -10^{-300}) for large s if some standard math library functions for the t distribution are used (e.g. `tpdf` and `tcdf` in Matlab v. 6 release 12).

To better distinguish which system should receive samples in a given stage ($\mathcal{OCBA}$, 0-1_1 and $\mathcal{LL}_1$), numerical stability was increased by evaluating the system that maximizes $\log(\text{EAPCS}_i - \text{APCS})$ (for $\mathcal{OCBA}$) and $\log \text{EVI}_{(i)}$ (for 0-1_1 and $\mathcal{LL}_1$). In particular, for $\mathcal{OCBA}$, set $p_j = \Phi_{\nu_{(j)(k)}}(d^*_{jk})$, $\tilde{p}_j = \Phi_{\tilde{\nu}_{(i)(k)}}(\tilde{\lambda}_{ik}^{1/2} d_{(i)(k)})$ and $\sum_j = \sum_{j:(j)\neq(k)}$ for convenience. Then

$\log(\text{EAPCS}_i - \text{APCS})$

$$= \begin{cases} \text{for } i \neq (k): \quad \sum_j \log(1-p_j) - \log(1-p_i) + \log p_i + \\ \qquad\qquad + \log\left[-\left(\exp(\log \tilde{p}_i - p_i) - 1\right)\right] \\ \text{for } i = (k): \quad \sum_j \log(1-\tilde{p}_j) + \\ + \log\left[-\left(\exp\left(\sum_j \log(1-p_j) - \sum_j \log(1-\tilde{p}_j)\right) - 1\right)\right] \end{cases} \tag{2.22}$$

These transformations are useful, because $\log(1+x) =$ `log1p(x)` and $\exp(x) - 1 =$ `expm1(x)` have increased accuracy for x near 0. In rare cases, we computed $\text{EAPCS}_i < \text{APCS}$, which we handled by setting $\log(\text{EAPCS}_i - \text{APCS})$ to $-\infty$.

For calculating $\log \text{EVI}$, we need $\log \Phi_\nu(t)$ and $\log \Psi_\nu(t)$. If the numerical stability does not suffice to calculate $\log \Phi_\nu(t)$ (underflow error) we derive bounds for $\log \Phi_\nu(t)$ based on the following property of the cdf of a t-distribution ([Evans, Hastings, and Peacock 1993]),

$$\Phi_\nu(t) = \begin{cases} \frac{1}{2}\beta^{\text{inc}}_{\text{reg}}(\frac{\nu}{2}, \frac{1}{2}, \frac{\nu}{\nu+t^2}) & \text{if } t \leq 0 \\ 1 - \frac{1}{2}\beta^{\text{inc}}_{\text{reg}}(\frac{\nu}{2}, \frac{1}{2}, \frac{\nu}{\nu+t^2}) & \text{if } t > 0, \end{cases} \tag{2.23}$$

where $\beta^{\text{inc}}_{\text{reg}}(a,b,x) = \beta(a,b)^{-1} \int_0^x u^{a-1}(1-u)^{b-1} du$ is the incomplete beta function, and $\beta(a,b) = \Gamma(a)\Gamma(b)/\Gamma(a+b)$ denotes the beta function. A

lower bound for $\log \Phi_\nu(-t)$ for $t > 0$ can be derived as follows. If $f(u) = u^{a-1}(1-u)^{b-1}$, then $f(0) = 0$, $f'(u) \geq 0$ and $f''(u) > 0$ for $a = \frac{\nu}{2} > 1$, $b = \frac{1}{2}$ and all $u \in [0,1]$. So the area below $f(u)$ over $[0,x]$ is always larger than the area below the tangent at $(x,\ f(x))$.

$$\begin{aligned} \log \Phi_\nu(-t) \quad &\geq \quad \tfrac{\nu}{2} \log \tfrac{\nu}{t^2+\nu} + \tfrac{1}{2} \log(1 - \tfrac{\nu}{t^2+\nu}) \\ &\qquad - \log\left((\tfrac{\nu}{2} - 1)(1 - \tfrac{\nu}{t^2+\nu}) + \tfrac{1}{2}\tfrac{\nu}{t^2+\nu}\right) - \log 2 \end{aligned} \tag{2.24}$$

For the upper bound, recall Equation (2.9). As $\Psi_\nu(t) > 0$ for all $t > 0$ we obtain $\Phi_\nu(-t) < \frac{1}{t}\frac{\nu+t^2}{\nu-1}\phi_\nu(t)$, so

$$\log \Phi_\nu(-t) \quad < \quad \log \tfrac{\nu/t+t}{\nu-1} + \log \phi_\nu(t) \tag{2.25}$$

The bounds in Equation (2.24) and Equation (2.25) help for Φ_ν but cannot directly be used for bounds on $\Psi_\nu(t)$. But the numerical stability can be increased by bringing the log inside the calculation of $\log \Psi_\nu(t)$:

$$\begin{aligned} \log \Psi_\nu(t) \quad &= \quad \log \tfrac{t^2+\nu}{\nu-1} + \log \phi_\nu(t) \\ &\qquad + \log\left[1 - \tfrac{t(\nu-1)}{(t^2+\nu)} \exp(\log \Phi_\nu(-t) - \log \phi_\nu(t))\right] \end{aligned} \tag{2.26}$$

Allocating based upon $\log$ EVI rather than EVI for 0-1_1 and $\mathcal{LL}_1$ improved performance when extreme guarantees of correct selection are sought (α^* or β^* close to 0), but required 50% more CPU time to determine the allocation, on average.

Collisions, due to the EVI being not numerically unique because of the interval bounds in Equation (2.22) through Equation (2.26), occurred almost always with 0-1_1 (the lower bound is loose) if high evidence levels for correct selection were required. Collisions occurred often with $\mathcal{LL}_1$. In the numerical experiments for $\mathcal{OCBA}$, 0-1_1 and $\mathcal{LL}_1$, if there was no clearly defined best (a collision), and the EVI or $\text{EAPCS}_i - \text{APCS}$ was not numerically different from 0 for *any* system (with interval arithmetic), then we repeatedly doubled τ for purposes of calculating EVI or $\text{EAPCS}_i - \text{APCS}$, until at least one system was numerically greater than 0. The 'winner' then received $\tau = 1$

replication. Usually, at most 3 doublings ($\tau = 8$) were sufficient to select a winner. If there was no clearly defined best because two or more systems whose EVI or $\text{EAPCS}_i - \text{APCS}$ had overlapping intervals but the intervals did not contain 0, then we allocated $\tau = 1$ replication to the system with the highest upper bound for the interval.

The performance of 0-1_1 and $\mathcal{LL}_1$ does not appear to be significantly hurt by the collisions, as the curves in the ($E[N]$, $\log(1 - \text{PCS}_{\text{iz}})$) and ($E[N]$, $\log \text{EOC}_{\text{iz}}$) planes appear relatively straight. Collisions occur rarely with $\mathcal{OCBA}$ and $\mathcal{OCBA}_{\text{LL}}$, but there is some slight bend to the right for low values of α^* or EOC bounds. That may suggest a potential inefficiency due to another numerical issue that we have not yet identified.

Procedures $\mathcal{KN}$++, $\mathcal{LL}$, and 0-1 did not experience numerical stability problems with collisions.

2.2.7 Summary of Tested Procedures

In addition to $\mathcal{KN}$++, we tested eight different allocation procedures, namely

- Equal, which allocates an equal number of samples to each alternative,
- two VIP procedures that allocate with a PCS (denoted 0-1) or EOC (denoted $\mathcal{LL}$) criterion,
- two corresponding small-sample EVI allocation (denoted 0-1_1 and $\mathcal{LL}_1$),
- three OCBA procedures that allocate with a PCS (denoted $\mathcal{OCBA}$), PGS (denoted $\mathcal{OCBA}_{\delta^*}$), and EOC (denoted $\mathcal{OCBA}_{\text{LL}}$) criterion.

Each allocation except for $\mathcal{KN}$++ was used in combination with each of three stopping rules defined in Section 2.2.1 ($\mathcal{S}$, $\text{PGS}_{\text{Slep},\delta^*}$, and EOC_{Bonf}). Overall, this resulted in 25 different procedures. So many variations were tested (a) to be inclusive and match all combinations in order to better understand the relative influence of each, (b) to unify separate streams of literature where small numbers of variants are compared at a time and numerical tests do not tend to be comparable, and (c) show the improvement

in both VIP and OCBA procedures with stopping rules other than $\mathcal{S}$ (the default in all past VIP and OCBA work).

We also tested the effect of including prior information about the means and variances in the VIP and OCBA configurations, as discussed in Section 2.4 below.

2.3 Evaluation Criteria

There are several ways to evaluate selection procedures, including the theoretical, empirical, and practical perspectives. Section 2.2 indicates that the three approaches make different basic assumptions, and that the most efficient representatives of each approach each make approximations of one sort or another. Theory that directly relates the different approaches is therefore difficult to develop.

We turn to the empirical and practical perspectives. The efficiency of a procedure is a frequentist measure of evidence for correct selection ($\mathrm{PCS}_{\mathrm{iz}}$, $\mathrm{PGS}_{\mathrm{iz},\delta^*}$ and $\mathrm{EOC}_{\mathrm{iz}}$) as a function of the average number of replications $E[N]$. As a function of each problem instance and sampling allocation, the stopping rule parameters *implicitly* define *efficiency curves* in the $(E[N], \mathrm{PCS}_{\mathrm{iz}})$ plane. [Dai 1996] proved exponential convergence for ordinal comparisons in certain conditions, so efficiency curves might be anticipated to be roughly linear on a logarithmic scale, $(E[N], \log(1 - \mathrm{PCS}_{\mathrm{iz}}))$, where $1 - \mathrm{PCS}_{\mathrm{iz}} = \mathrm{PICS}_{\mathrm{iz}}$. Efficiency curves for $\mathrm{EOC}_{\mathrm{iz}}$ and $\mathrm{PGS}_{\mathrm{iz},\delta^*}$ are defined similarly. 'More efficient' procedures have curves that are below those of other procedures.

Efficiency curves ignore the question of how to set a procedure's parameters to achieve a particular $\mathrm{PCS}_{\mathrm{iz}}$ or $\mathrm{EOC}_{\mathrm{iz}}$. As a practical matter, one expects some deviation between a stopping rule target, say $\mathrm{PCS} \geq 1 - \alpha^*$, and the actual $\mathrm{PCS}_{\mathrm{iz}}$ achieved. The deviation between the desired and realized performance is measured with *target curves* that plot $(\log \alpha^*, \log(1 - \mathrm{PCS}_{\mathrm{iz}}))$ for PCS-based targets $1 - \alpha^*$, and $(\log \beta^*, \log \mathrm{EOC}_{\mathrm{iz}})$ for opportunity cost targets β^*. Procedures whose target curves follow the diagonal $y = x$ over a range of problems are 'controllable' in that it is possible to set parameter values to obtain a desired level of correct selection. 'Conservative' procedures

have target curves that tend to be below $y = x$, and are said to 'overdeliver' because the frequentist measure for correct selection exceeds the desired target. An IZ procedure may therefore not be controllable if it is conservative in an unpredictable way. A controllable procedure may not have a PCS_{iz} guarantee if the target curve goes slight above or below the diagonal.

2.4 Test Bed Structure

A large number of problem instances assessed the strengths and weaknesses of each procedure. We varied the number of systems, the first stage sampling size, and the configuration of the means and variances. We tested random problem instances and the ability to use prior information about the unknown means.

Classic Problem Instances In a **slippage configuration (SC)**, the means of all systems except the best are tied for second best. A SC is identified by the number of systems, the difference in means of the best and each other system, and the variances of each system. The parameters δ, ρ describe the configurations we tested.

$$\begin{aligned} X_{1j} &\overset{iid}{\sim} \texttt{Normal}\left(0, \sigma_1^2\right) \\ X_{ij} &\overset{iid}{\sim} \texttt{Normal}\left(-\delta, \sigma_1^2/\rho\right) \text{ for systems } i = 2, \ldots, k \end{aligned}$$

If $\rho = 1$, then all systems have the same variance, and $\rho < 1$ means that the best system has a smaller variance. We set $\sigma_1^2 = 2\rho/(1+\rho)$ so that $\text{Var}[X_{1j} - X_{ij}]$ is constant for all $\rho > 0$.

In a **monotone decreasing means (MDM)** configuration, the means of all systems are equally spaced out, so that some systems are quite a bit inferior to the best. The parameters δ, ρ describe the configurations that we tested. The outputs were jointly independent, and we set σ_1^2 like in SC.

$$\begin{aligned} X_{1j} &\sim \texttt{Normal}\left(0, \sigma_1^2\right) \\ X_{ij} &\sim \texttt{Normal}\left(-(i-1)\delta, \sigma_1^2/\rho^{i-1}\right) \text{ for systems } i = 2, \ldots, k \end{aligned}$$

Values of $\rho < 1$ mean that better systems have a smaller variance. For sufficiently small ρ, the probability that the worst system has the best observed mean is higher than the probability for the second best system.

For the SC and MDM configurations we tested many (hundreds), but not all of the following parameter combinations: $k \in \{2, 5, 10, 20, 50\}$, $\delta \in \{0.25, 0.354, 0.5, 0.707, 1\}$, and $\rho \in \{0.125, 0.177, 0.25, 0.354, 0.5, 0.707, 1, 1.414, 2, 2.828, 4\}$, with $n_0 \in \{4, 6, 10\}$. We tested $B \in \{kn_0, \ldots, 2k\left(\frac{1}{\delta}\Phi^{-1}\left(\frac{\alpha^*}{k}\right)\right)^2\}$ for the budget stopping rule ($\mathcal{S}$), varied the indifference zone parameter relative to the difference in means $\delta^* \in \{0, 0.05, 0.1, \ldots, 0.6\}$ and $\alpha^* \in [0.001, 0.5]$ for PGS-based stopping rules ($\mathrm{PGS}_{\mathrm{Slep},\delta^*}$) and $\mathcal{KN}$++, and varied $\beta^* \in [0.001, 0.5]$ for ($\mathrm{EOC}_{\mathrm{Bonf}}$). For some configurations, we varied $\alpha^* \in [0.001, 1/k]$ for $\mathcal{KN}$++ as the achieved PICS for $\alpha^* = 0.5$ was below 0.01.

Random Problem Instances Assessments of selection procedures in the literature usually apply procedures to a specific set of structured problems, as above. A **Random problem instance (RPI)** may be more realistic in the sense that problems faced in practice are typically not in the SC or MDM configuration. In each RPI experiment below, the output is again jointly independent, $X_{ij} \overset{iid}{\sim} \texttt{Normal}\,(\mu_i, \sigma_i^2)$ for $i = 1, \ldots, k$, conditional on the problem instance. The problem instance $\boldsymbol{\chi}$ is sampled randomly prior to applying a selection procedure. Correct selection metrics are generalized to be expectations over the sampling distribution, e.g. $\mathrm{PCS}_{\mathrm{iz}} = E_{\boldsymbol{\chi}}[\mathrm{PCS}_{\mathrm{iz}}(\boldsymbol{\chi})]$.

The first RPI experiment (**RPI1**) samples $\boldsymbol{\chi}$ from the normal-inverse gamma family. A random $\boldsymbol{\chi}$ is generated by sampling the σ_i^2 independently, then sampling the M_i, given σ_i^2,

$$\begin{aligned} p(\sigma_i^2) &\sim \texttt{InvGamma}\,(\alpha, \beta) \\ p(M_i \,|\, \sigma_i^2) &\sim \texttt{Normal}\left(\mu_0, \sigma_i^2/\eta\right). \end{aligned} \tag{2.27}$$

If $S \sim \texttt{InvGamma}\,(\alpha, \beta)$, then $E[S] = \beta/(\alpha-1)$, $S^{-1} \sim \texttt{Gamma}\,(\alpha, \beta)$, $E[S^{-1}] = \alpha\beta^{-1}$ and $\mathrm{Var}[S^{-1}] = \alpha\beta^{-2}$. Increasing η makes the means more similar and therefore the problem harder. We set $\beta = \alpha - 1 > 0$ to standardize

the mean of the variance to be 1, and set $\mu_0 = 0$. Increasing α reduces the difference between the variances. We tested many combinations out of $k \in \{2, 5, 10, 20\}$, $\eta \in \{0.0625, 0.0884, 0.125, \ldots, 4\}$, $\alpha \in \{2.5, 100\}$, and $n_0 \in \{4, 6, 10\}$. The derivations of the VIP and OCBA procedures correspond to $\eta \to 0$.

The RPI1 experiment permits a test of whether the VIP and OCBA procedures can benefit from using the sampling distribution of $\boldsymbol{\chi}$ in Equation (2.27) to describe prior judgement about the means and variances of each system. Section 2.2 does not allow for this directly, but the mathematical development to do so was provided elsewhere for the VIP ([Chick and Inoue 1998; Chick and Inoue 2001b]). In summary, the posterior distribution of M_i, given the prior distribution in Equation (2.27) and data $\mathcal{E}_i = (x_{i1}, \ldots, x_{in_i})$, is

$$\begin{aligned} p(\sigma_i^2 \,|\, \mathcal{E}_i) &\sim \texttt{InvGamma}\,(\alpha', \beta'), \\ p(M_i \,|\, \sigma_i^2, \mathcal{E}_i) &\sim \texttt{Normal}\left(\mu_0', \sigma_i^2/\eta'\right) \end{aligned}$$

where $\alpha' = \alpha + n_i/2$, $\beta' = \beta + (\frac{\eta n_i}{\eta + n_i}(\mu_0 - \bar{x}_i)^2 + \sum_{j=1}^{n_i}(x_{ij} - \bar{x}_i)^2)/2$, $\mu_0' = \frac{\eta\mu_0 + n_i\bar{x}_i}{\eta + n_i}$, and $\eta' = \eta + n_i$. To apply that result to all VIP procedures in Section 2.2.3, substitute each $\bar{x}_i$ with μ_0'; replace each s_i^2 with β'/α'; and replace each n_i with η', except in the degrees of freedom, where $n_i - 1$ should be replaced with $2\alpha'$. The $\mathcal{OCBA}$ has previously always assumed a noninformative prior distribution. Analogous substitutions allow $\mathcal{OCBA}$ and $\mathcal{OCBA}_{\text{LL}}$ to use other prior distributions for the unknown means and variances.

A second RPI experiment (**RPI2**) samples problem instances from a distribution other than normal-inverted gamma to remove any potential advantage for the VIP and OCBA approaches. There is no objectively best distribution for sampling problem instances. We chose RPI2 to independently sample from:

$$\begin{aligned} \sigma_i^2 &\sim \texttt{InvGamma}\,(\alpha, \beta) \\ M_i \,|\, \sigma_i^2 &\sim (-1)^a \texttt{Exponential}\left((\eta/\sigma_i^2)^{1/2}\right), \end{aligned}$$

where the mean of an `Exponential` (λ) distribution is $1/\lambda$. $-$`Exponential` (λ) is a notational abbreviation for a distribution, where the realizations of an exponential are inverted. There are typically several competitors for the best if $a = 1$ and few competitors for the best if $a = 0$. A larger η makes for harder problems with closer means. Heterogeneity in the variances is controlled with α and β. We tested the same parameters as in RPI1.

Summary of Configurations The SC favors IZ procedures in that IZ procedures provide a minimal target performance with respect to a least favorable configuration (LFC), and for many IZ procedures an SC with $\delta = \delta^*$ is a LFC. The RPI1 (η near 0) may favor the VIP and OCBA, as the derivation of those procedures assume prior probability models that are similar to the sampling distribution of the problem instances. The MDM, RPI1 (larger η) and RPI2 experiments do not favor any procedure in this chapter.

2.5 Empirical Results

This section summarizes the qualitative features of the analysis. Plots were generated with 10^5 macroreplications for each combination of problem instance, sampling allocation, and stopping rule parameter value. To improve the significance of contrasts between different procedures, common random numbers (CRN) were used to generate common configurations for RPI experiments, and to synchronize the samples observed across procedures. CRN were *not* used between systems. $\mathcal{KN}{+}{+}_{\delta^*}$ refers to $\mathcal{KN}{+}{+}$ with the given indifference zone parameter. By default, $n_0 = 6$ unless specified otherwise.

Two Systems, SC/MDM. When there are $k = 2$ systems, the SC and MDM configurations are equivalent, and PICS_{iz} is proportional to EOC_{iz}. When the variances are equal, it is optimal to sample equally often from both systems (e.g. [Gupta and Miescke 1994]), so Procedure Equal samples optimally. Figure 2.3 demonstrates the effect of different stopping rules on efficiency (similar effect for different δ). The EOC-based stopping rule is more efficient than the PCS-based stopping rule. Both are much more efficient than

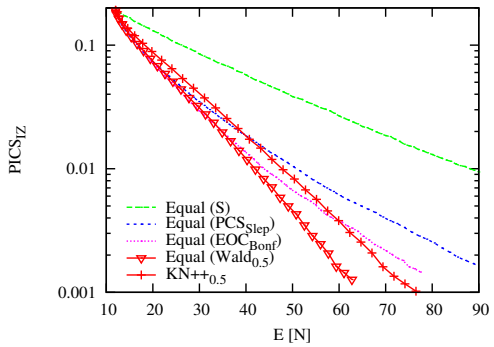

Figure 2.3: Efficiency of Equal allocation with different stopping rules and $\mathcal{KN}$++ with $\delta^* = \delta$ (SC, $k = 2$, $\delta = 0.5$, $\rho = 1$).

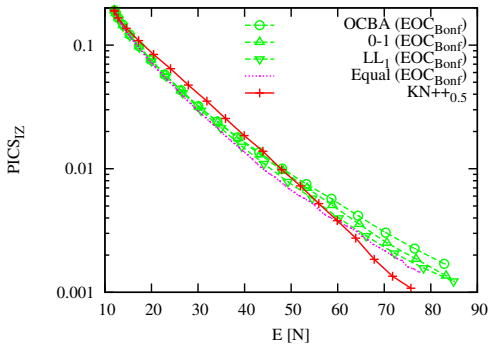

Figure 2.4: Efficiency of different allocation procedures and $\mathcal{KN}$++ with $\delta^* = \delta$ (SC, $k = 2$, $\delta = 0.5$, $\rho = 1$).

stopping after a fixed budget ($\mathcal{S}$) because any additional sampling is adapted to the level of evidence observed so far.

If $k = 2$, then $\mathcal{KN}$++ samples each system equally often until the stopping criterion is met, so it is equivalent to the Equal allocation with a special stopping rule. For higher PICS, Equal(EOC_{Bonf}) in this case is more efficient than $\mathcal{KN}$++, while for very low PICS, $\mathcal{KN}$++ beats Equal(EOC_{Bonf}). Figure 2.3 also shows that the efficiency curve for $\mathcal{KN}$++ and Equal allocation with the $\mathcal{S}$ and Wald stopping rule is straighter than for the PCS or EOC stopping rules. The Bayesian stopping rules cause a slight curvature.

Allocating equally and stopping with Wald's SPRT is the most efficient procedure. It outperforms all other procedures for each stopping value.

Figure 2.4 shows that of all allocations with the EOC_{Bonf} stopping rule, Equal performs most efficiently (it is optimal for this particular setting), with $\mathcal{LL}_1$, 0-1 and $\mathcal{OCBA}$ following. $\mathcal{LL}$ performs identical to 0-1, and $0\text{-}1_1$ is very similar to 0-1 for this problem (not shown). A similar precedence is observed for the PCS_{Slep} and $\text{PGS}_{\text{Slep},\delta^*}$ stopping rules. For the $\mathcal{S}$ stopping rule, all VIP and OCBA allocations perform about the same as Equal (not shown). The relative ordering of the stopping rules for the equal allocation, namely EOC_{Bonf} beats PCS_{Slep} which beats $\mathcal{S}$, is also observed for *all* VIP and OCBA allocations with a similar order of magnitude difference (not shown).

With adaptive stopping rules (EOC_{Bonf}, PCS_{Slep}, $\text{PGS}_{\text{Slep},\delta^*}$), a large

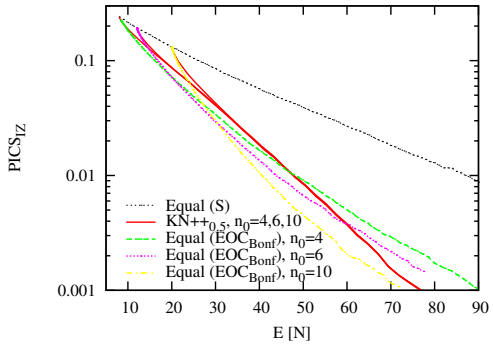

Figure 2.5: Influence of n_0 (SC, $k = 2$, $\delta = 0.5$, $\rho = 1$).

number of initial samples per system, n_0, limits the opportunity to make an early selection, but a small n_0 increases the probability of poor estimates of the output mean and variance. For the posterior marginal distributions of the unknown means to have a finite variance, we require $n_0 \geq 4$ (see text after Equation (2.2)). Figure 2.5 shows that increasing n_0 in Procedure Equal(EOC_{Bonf}) increases the number of samples required to reach relatively low levels of evidence for correct selection, but increases the efficiency of the procedure to reach high levels of evidence for correct selection. The differences in the curves are predominantly attributed to output that causes sampling to stop after very few samples, due to misleadingly low variance and PICS estimates. The OCBA and VIP procedures behave similar to Equal in this respect for each stopping rule. With the nonadaptive stopping rule ($\mathcal{S}$), they seem insensitive to n_0. Procedure $\mathcal{KN}$++ is quite insensitive to n_0 (regardless of k, not shown).

The tests above use $\delta^* = \delta$ for $\mathcal{KN}$++ and $\delta^* = 0$ for $\text{PGS}_{\text{Slep},\delta^*}$. That choice seems natural for $\mathcal{KN}$++, since $\delta^* = \delta$ is a LFC for many IZ procedures. For $\text{PGS}_{\text{Slep},\delta^*}$, the choice $\delta^* = 0$ seems the natural choice for PCS_{iz} efficiency. But the parameter δ^* has a strong effect on efficiency and target curves. Setting δ^* to obtain a desired PICS_{iz} seems challenging (an observation for most configurations tested, not just SC).

Figure 2.6 shows the influence of δ^* on PCS_{iz} efficiency (the mean number of samples required to obtain a specified level PICS_{iz}). For small PICS_{iz}, there may exist settings for δ^* so that $\mathcal{KN}$++ and Equal($\text{PGS}_{\text{Slep},\delta^*}$) are

more efficient than Equal(EOC_{Bonf}). To see this, note that when $\text{PICS}_{\text{iz}} = 0.005$ and 0.01, the curves for $\mathcal{KN}$++ and Equal($\text{PGS}_{\text{Slep},\delta^*}$) go below the horizontal lines. The horizontal lines show mean number of samples required by Equal(EOC_{Bonf}) to reach the corresponding PICS_{iz} level. The value of δ^* that is needed to obtain the minimal mean number of samples depends on the problem instance. The curvature is induced by two concurring effects: For a given level of PICS_{iz}, α^* has to be decreased if δ^* is increased. The mean number of samples are *increased* by decreasing α^* but are also *decreased* by increasing δ^*. Both effects are nonlinear, which gives the curvature in Figure 2.6.

Since the problem instance is unknown in practice, it is not clear how to set δ^* in general. Figure 2.7 illustrates that δ^* also has a strong effect on the target performance. There is no obvious way to set δ^*, α^* to reliably achieve a given PICS_{iz} goal. Choosing a small δ^* may result in sampling way beyond what is actually needed, in particular for $\mathcal{KN}$++ and Wald. While $\delta^* = \delta$ would yield very good target performance for $\mathcal{KN}$++, δ is usually unknown. Equal($\text{PGS}_{\text{Slep},\delta^*}$) significantly underdelivers for $\delta^* \geq 0.3$, i.e. values smaller than the true difference between the best and second best system.

The EOC_{Bonf} stopping rule is quite sensitive to the difference between the two systems, δ (see Figure 2.8). It slightly underdelivers EOC_{iz} for small δ, and significantly overdelivers for large δ. A similar behavior is observed for $\mathcal{KN}$++ here when control for EOC_{iz} is attempted with $\beta^* = \delta^*\alpha^*$.

Overall, for SC with $k = 2$, Equal(EOC_{Bonf}) and $\mathcal{KN}$++ seem the most efficient. No procedure is fully controllable for SC, $k = 2$. The remarks so far presume a common variance ($\rho = 1$). When $\rho \neq 1$, the equal allocation is not optimal, and $\mathcal{KN}$++ and Equal allocation are slightly less efficient (not shown).

SC with $k > 2$ systems. It is not optimal to sample each system equally often if $k > 2$. Figure 2.9 shows the efficiency of different allocation rules for the EOC_{Bonf} stopping rule and $k = 10$ systems. The settings are comparable to Figure 2.4 with $k = 2$ systems. While Equal is optimal for $k = 2$, it performs worst for $k = 10$. The most efficient allocations are $\mathcal{OCBA}_{\text{LL}}$

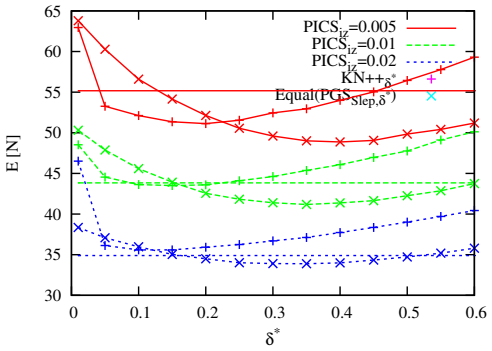

Figure 2.6: Influence of δ^* on mean number of samples to obtain a desired PICS_{IZ}, for $\mathcal{KN}$++ and Equal($\text{PGS}_{\text{Slep},\delta^*}$) (SC, $k = 2$, $\delta = 0.5$, $\rho = 1$).

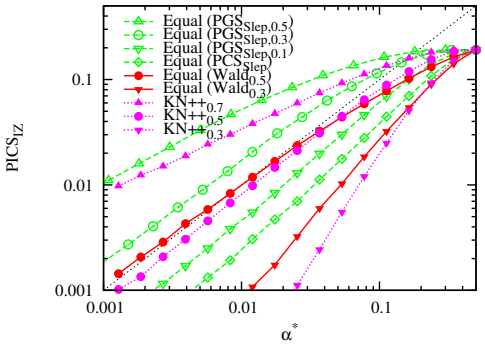

Figure 2.7: Influence of δ^* on $\mathcal{KN}$++ and Equal($\text{PGS}_{\text{Slep},\delta^*}$) *target* performance. (SC, $k = 2$, $\delta = 0.5$, $\rho = 1$).

and $\mathcal{OCBA}$, then $\mathcal{LL}$ closely behind. $\mathcal{KN}$++ is much less efficient than $\mathcal{OCBA}_{\text{LL}}$, $\mathcal{OCBA}$ and $\mathcal{LL}$ (each with EOC_{Bonf} as a stopping rule). The difference between the efficiency of the Bayesian procedures relative to the Equal and $\mathcal{KN}$++ procedures increases with k (tested $k = 2, 5, 10, 20$). The qualitative nature of the claim does not change as δ and ρ are individually varied from the values used for the plot.

Procedures $\mathcal{LL}_1$ and 0-1_1 would perform similarly to $\mathcal{KN}$++ in Figure 2.9 (not shown). In fact, the small sample procedures turned out to be generally less efficient than their counterparts $\mathcal{LL}$ and 0-1. The implication is that the use of the Bonferroni and Welch approximations by $\mathcal{LL}$ and 0-1 causes less deviation from optimal sampling than myopically allocating samples one at a time to reduce EVI in $\mathcal{LL}_1$ and 0-1_1.

Other observations for $k = 2$ also hold for $k = 5, 10$ and 20, including: the precedence of the effectiveness of stopping rules (EOC_{Bonf} beats PCS_{Slep} which beats $\mathcal{S}$); the importance of a sufficiently large n_0 for the VIP and OCBA procedures; the sensitivity of $\mathcal{KN}$++ to δ^* but not n_0.

Monotone Decreasing Means. For MDM and $k > 2$ equal allocation is no longer optimal. Figure 2.10 shows the optimal allocation compared to $\mathcal{OCBA}$ and $\mathcal{LL}$. The curves for the loss-based procedures ($\mathcal{OCBA}_{\text{LL}}$ not

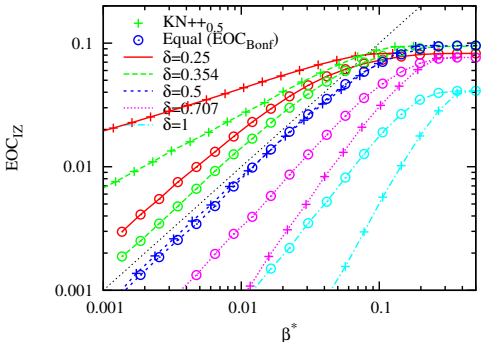

Figure 2.8: Influence of δ on target performance (SC, $k = 2$, $\rho = 1$).

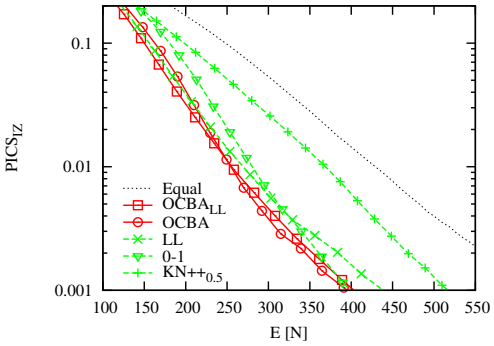

Figure 2.9: Efficiency with $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rule (SC, $k = 10$, $\delta = 0.5$, $\rho = 1$).

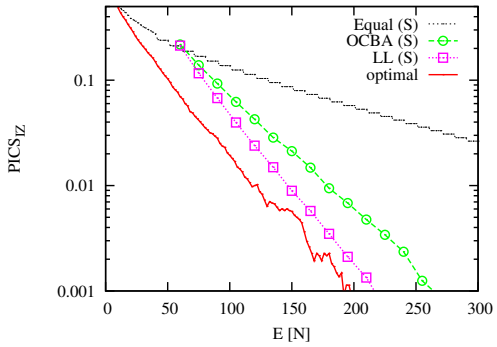

Figure 2.10: Comparison of OCBA with the optimal allocation rule. The line for optimal allocation is more rugged due to fewer simulations (10^4 instead of 10^5); MDM ($k = 10$, $\delta = 0.5$, $\rho = 1$).

shown) are nearly parallel to the optimal allocation, which indicates that these procedures allocate almost optimal after the initial n_0 samples.

MDM adds the complication that $\mathrm{EOC}_{\mathrm{iz}}$ is not proportional to $\mathrm{PCS}_{\mathrm{iz}}$ when $k > 2$. Figure 2.11 illustrates that for the Equal allocation, the $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rule outperforms $\mathrm{PCS}_{\mathrm{Slep}}$, which beats $\mathcal{S}$, an order observed for *all* VIP and OCBA allocations, and *all* MDM configurations tested. Notably, the $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rule outperforms $\mathrm{PCS}_{\mathrm{Slep}}$ not only for $\mathrm{EOC}_{\mathrm{iz}}$ efficiency, but also for $\mathrm{PCS}_{\mathrm{iz}}$ efficiency. As for SC, the $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$ stopping rule can be more efficient than the $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rule for some (but not all) δ^*.

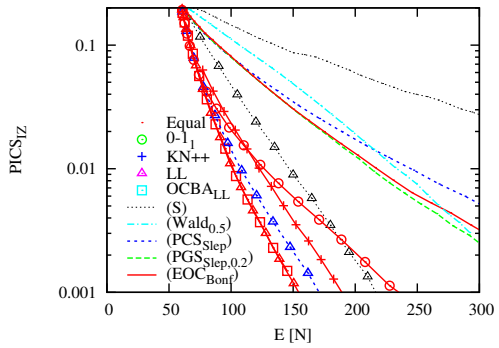

Figure 2.11: Different stopping rules (line types) and allocation procedures (symbols) (MDM, $k = 10$, $\delta = 0.5$, $\rho = 1$).

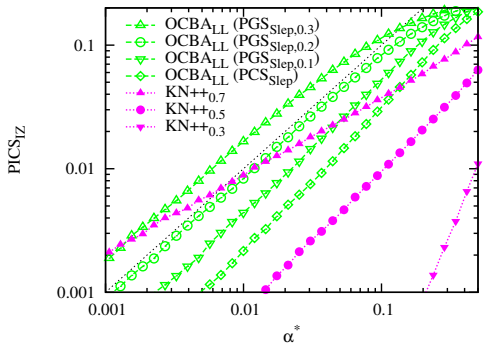

Figure 2.12: Target graphs for $\mathcal{KN}$++ and 0-1 (MDM, $k = 10$, $\delta = 0.5$, $\rho = 1$).

Figure 2.11 also exemplifies an observation for MDM runs that $\mathcal{KN}$++ with $\delta^* = \delta$ is typically more efficient than the original VIP and OCBA procedures, which used the $\mathcal{S}$ stopping rule. However, the original VIP and OCBA allocations with the new EOC_{Bonf} stopping rule are more efficient than $\mathcal{KN}$++. Although stopping with Wald was the most efficient for $k = 2$, it is the worst stopping rule in MDM. This holds for Wald used in the allocation, too (not shown). The most efficient allocations are $\mathcal{LL}$ and $\mathcal{OCBA}_{\text{LL}}$ (along with 0-1 and $\mathcal{OCBA}$, not shown). The small sample procedures, $0\text{-}1_1$ and $\mathcal{LL}_1$, were only competitive for high PICS values.

For MDM, the EOC-based allocations typically outperform the corresponding 0-1-based allocations for all stopping rules, the difference being particularly sizable for the $\mathcal{S}$ stopping rule (not shown). This is true not only for efficiency with respect to EOC_{iz}, but also to PCS_{iz}. For each MDM configuration tested, $\mathcal{LL}(\text{EOC}_{\text{Bonf}})$ and $\mathcal{OCBA}_{\text{LL}}(\text{EOC}_{\text{Bonf}})$ are not statistically different for efficiency. Those two procedures also perform roughly similar in target plots for most configurations and follow the target diagonal reasonably well in most cases, although some differences with no discernable pattern were observed in some target plots.

Figure 2.12 illustrates that $\mathcal{KN}$++ is again very sensitive to the parameter δ^* (difficult to control). While $\mathcal{KN}$++ adheres to the target quite well

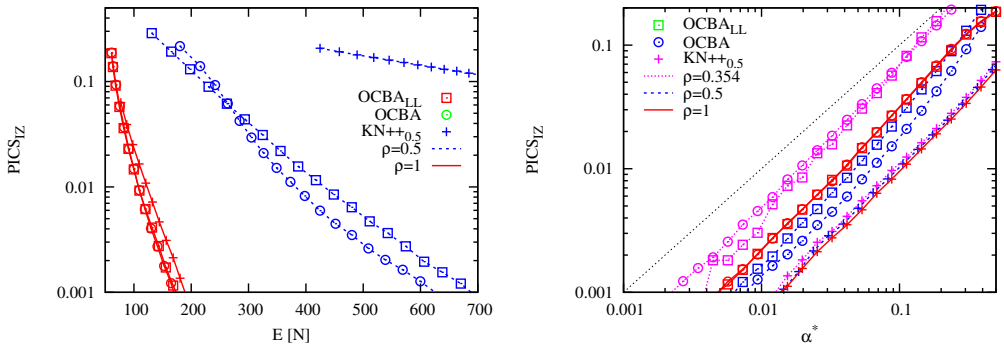

Figure 2.13: Effect of variance ratio ρ on efficiency (left panel) and target performance (right panel). (MDM, $k = 10$, $\delta = 0.5$). Procedures $\mathcal{OCBA}$ and $\mathcal{OCBA}_{\text{LL}}$ use PCS_{Slep} stopping rule. Procedures $\mathcal{OCBA}_{\text{LL}}$ and $\mathcal{LL}$ perform the same.

for SC when $\delta^* = \delta$, it significantly overdelivers even for this setting for the MDM configuration. Procedure $\mathcal{OCBA}_{\text{LL}}(\text{PGS}_{\text{Slep},\delta^*})$ is also sensitive to the parameter δ^*. While the target curves for $\mathcal{OCBA}_{\text{LL}}(\text{PGS}_{\text{Slep},\delta^*})$ shift roughly parallel to the diagonal, the target curves of $\mathcal{KN}{+}{+}$ change in slope. As for SC, $\mathcal{KN}{+}{+}$ and $\text{PGS}_{\text{Slep},\delta^*}$ are quite sensitive to δ and thus not controllable.

Figure 2.13 illustrates the effect of the output variance ratio ρ on different procedures for MDM with $k = 10$. With $\rho = 1$ (equal variance), the best PCS_{Slep} procedures perform somewhat more efficiently than $\mathcal{KN}{+}{+}$. Increasing ρ (best systems have larger variance) has little effect on the relative performance of the procedures. Decreasing ρ to 0.5 (very large variance for the worst systems) increases the total number of samples for all procedures, but particularly deteriorates the efficiency of $\mathcal{KN}{+}{+}$ (in far right of efficiency plot) relative to the PCS_{Slep} procedures. Procedure $\mathcal{KN}{+}{+}$ also overdelivers more than some other procedures (right panel). The target curves for all procedures are relatively insensitive to ρ (the target relies primarily on the difference between the two systems competing most for best, so efficiency is affected more than the target curve).

Procedures $\mathcal{OCBA}_{\text{LL}}$ and $\mathcal{LL}$ are somewhat more efficient than $\mathcal{OCBA}$ and 0-1 at high PICS levels, but the reverse may be true in many MDM and SC configurations for very low PICS levels (not shown).

For all SC and MDM configurations and almost all sampling procedures,

the efficiency curves exhibit a certain curvature. We found several explanations for curved efficiency lines for the OCBA and VIP procedures. One, a small n_0 leads to poorer variance estimates initially, with a potential for either (a) early stopping if PICS is strongly underestimated, or (b) a massive number of samples being required if an extremely low PICS or EOC is desired, initial estimates suggest that the best system is worst, and the procedure then tries to distinguish between the equal systems in the SC. Both cases are alleviated by increasing n_0. Two, the test bed pushed the procedures to new limits for numerical stability. Preliminary efficiency plots for some procedures were somewhat more curved than those presented here. Section 2.2.6 describes computational techniques to reduce that curvature. We believe that this cause was eliminated. Three, exponential convergence results for ordinal comparisons are asymptotic and available for only some procedures, so straight lines might not be expected at all levels of PICS_{iz} and EOC_{iz} for all procedures.

Random Problem Instances 1. For the RPI experiments, it is necessary to choose $\delta^* > 0$ for the $\text{PGS}_{\text{Slep},\delta^*}$ stopping rule because there is a reasonable probability that the two best systems have very similar means, in which case $\delta^* = 0$ results in excessive sampling. Therefore $\delta^* = 0$ (or PCS_{Slep}) is to be avoided in practice. We examine efficiency for the probability of a bad selection, $\text{PBS}_{\text{iz},\delta^*}$, instead of PICS_{iz}, in this section.

For basically *all* RPI settings, the $\mathcal{LL}$, $\mathcal{OCBA}_{\text{LL}}$ and $\mathcal{OCBA}_{\delta^*}$ allocation rules are more or less equally efficient. The 0-1 allocation is generally less efficient (it is derived with more approximations, and wastes samples trying to distinguish between two very close competitors in the RPI) and Equal is worst. See Figure 2.14 for the $\mathcal{S}$ stopping rule (which may be needed if an analysis has a strict time constraint).

While the differences among the allocation rules are rather small for RPI, the performance of the stopping rules varies widely. Figure 2.15 compares different stopping rules in combination with Equal allocation based on $\text{PGS}_{\text{iz},\delta^*}$ efficiency. $\text{PGS}_{\text{Slep},\delta^*}$ stopping rule and 'matching' δ^* clearly performs best. By matching, we mean that the δ^* chosen for the stopping rule equals the

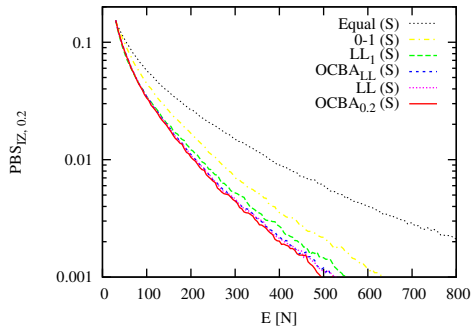

Figure 2.14: PBS_{δ^*} Efficiency of allocations with $\mathcal{S}$ stopping rule (RPI1, $k = 5$, $\eta = 1$, $\alpha = 100$).

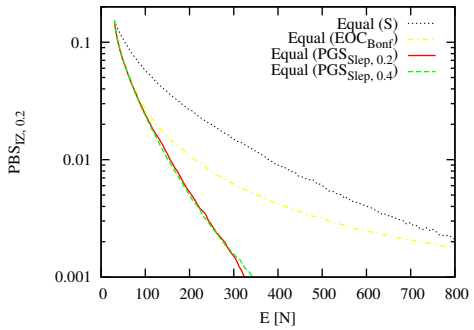

Figure 2.15: PBS_{δ^*} Efficiency of Equal allocation (RPI1, $k = 5$, $\eta = 1$, $\alpha = 100$).

value of δ^* chosen to measure $\mathrm{PGS}_{\mathrm{iz},\delta^*}$. While this is not necessarily the most efficient setting, it seems the most sensible. For RPI, the effect of δ^* on efficiency is relatively small (in Figure 2.15, $\mathrm{Equal}(\mathrm{PGS}_{\mathrm{Slep},0.2})$ and $\mathrm{Equal}(\mathrm{PGS}_{\mathrm{Slep},0.4})$ are almost identical), and the chosen matching setting yields a very good $\mathrm{PGS}_{\mathrm{iz},\delta^*}$ target performance.

For $\mathrm{EOC}_{\mathrm{iz}}$ efficiency, settings for δ^* exist so that $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$ is more efficient than the $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rule, see Figure 2.16 (left panel). Whether that finding is of practical relevance remains to be seen, as it is not yet clear how to set δ^* in $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$ to control $\mathrm{EOC}_{\mathrm{iz}}$ via $\beta^* = \delta^*\alpha^*$ (right panel).

Several efficiency curves, in particular $\mathrm{Equal}(S)$ in Figure 2.15 and Figure 2.16, are more curved for RPI1 than for the SC and MDM configurations. That curvature is largely due to a very large number of samples for a few very "hard" configurations (the best two systems have very close means and large variances).

Figure 2.17 compares three selection procedures with adaptive stopping rules, $\mathcal{KN}{+}{+}$, $\mathrm{Equal}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$, and $\mathcal{OCBA}_{\delta^*}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$. As is typical for the RPI1 problems tested, $\mathcal{OCBA}_{\delta^*}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$ outperforms $\mathcal{KN}{+}{+}$ for efficiency (left panel) and controllability (right panel). When moving from $\alpha = 100$ (very similar variances for each system) to $\alpha = 2.5$ (very different variances), the efficiency of $\mathcal{OCBA}_{\delta^*}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$ improves and the efficiency of $\mathcal{KN}{+}{+}$ is basically not affected.

As can be seen in Figure 2.18, the sensitivity with respect to η in the

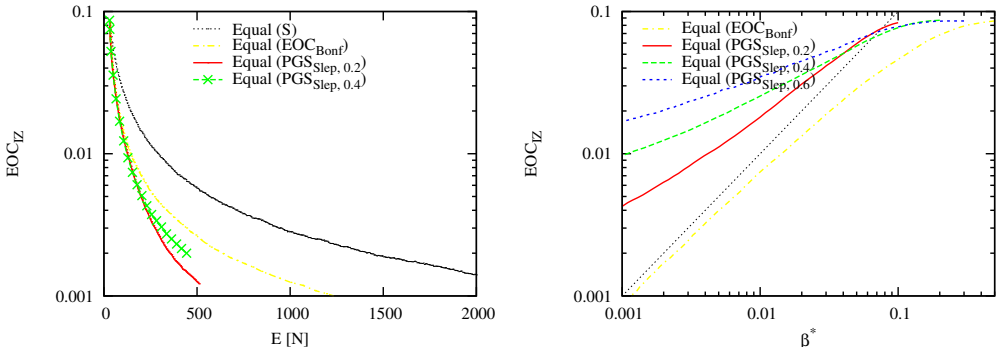

Figure 2.16: Efficiency (left) and target (right) for Equal allocation and different stopping rules w.r.t. $\mathrm{EOC}_{\mathrm{iz}}$ (RPI1, $k = 5$, $\eta = 1$, $\alpha = 100$). For $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$ stopping, β^* is approximated by $\alpha^*\delta^*$

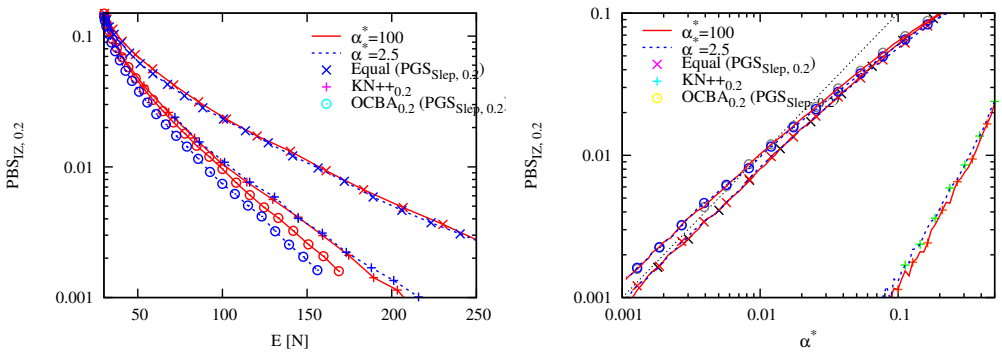

Figure 2.17: Efficiency and target for PCS-based procedures (RPI1, $k = 5$, $\eta = 1$, $\alpha = 2.5, 100$).

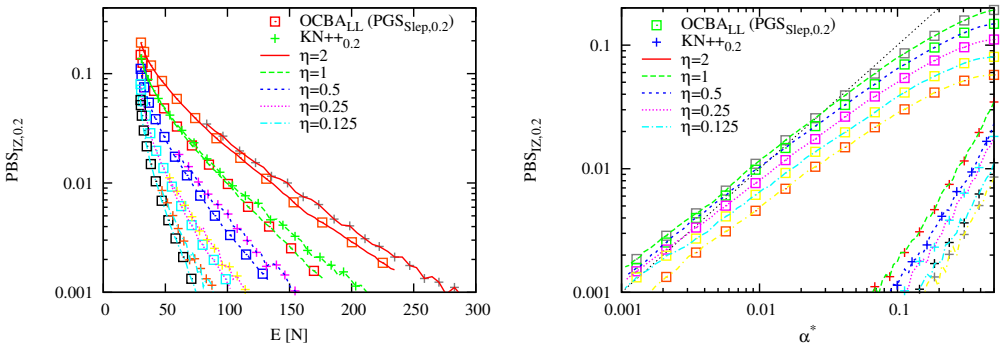

Figure 2.18: Influence of η on efficiency and target for PCS-based procedures (RPI1, $k = 5$, $\alpha = 100$).

RPI experiments is much smaller than the sensitivity with respect to δ observed for the SC and MDM configurations in Figure 2.8. Note that the difference between the best and second best system is proportional to η^{-2}. For efficiency (left panel) $\mathcal{OCBA}_{\mathrm{LL}}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$ slightly outperforms $\mathcal{KN}{+}{+}$. Regarding target performance, $\mathcal{KN}{+}{+}$ consistently and strongly overdelivers, while $\mathcal{OCBA}_{\mathrm{LL}}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$ meets the target rather well over all η tested. Also, procedures with the $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rule follow an $\mathrm{EOC}_{\mathrm{iz}}$ target well (not shown).

Observations from SC and MDM that apply to RPI1 as well include: The Bayesian procedures are generally quite sensitive to n_0, while $\mathcal{KN}{+}{+}$ is not; $\mathcal{KN}{+}{+}$ becomes less efficient relative to the Bayesian procedures as the number of systems k increases.

If prior knowledge on the distribution of means and variances is available, this can be integrated into the Bayesian procedures as described in Section 2.4. The benefit of doing so, when possible, is apparent in Figure 2.19 (left panel). The top line shows the efficiency of the standard procedure with Equal allocation and a budget stopping rule. This can be improved by switching to a flexible allocation procedure ($\mathcal{OCBA}_{\mathrm{LL}}(\mathcal{S})$), using an adaptive stopping rule ($\mathcal{OCBA}_{\mathrm{LL}}(\mathrm{EOC}_{\mathrm{Bonf}})$), and using prior information ($\mathcal{OCBA}^{\mathrm{prior}}_{\mathrm{LL}}(\mathrm{EOC}^{\mathrm{prior}}_{\mathrm{Bonf}})$). These changes reduce the mean number of samples required to achieve a loss of 0.01 from 291 to 164, to 94 and to 79. The target performance (right panel) is only slightly affected by incorporating prior

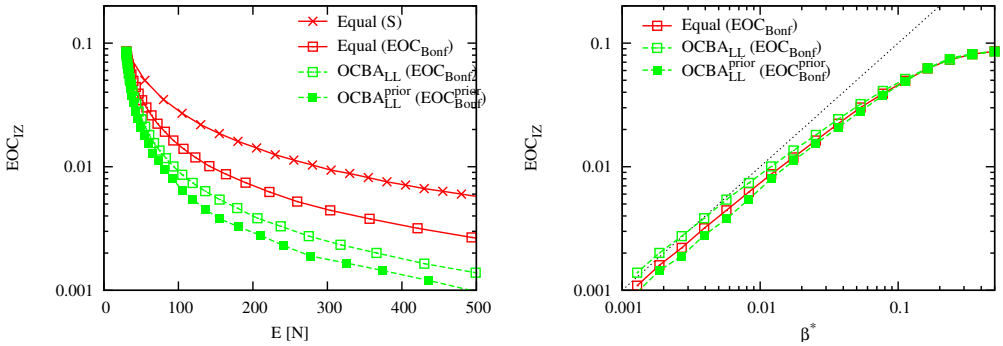

Figure 2.19: Effect of allocation, stopping rule and prior information (RPI, $k = 5$, $\eta = 1$, $\alpha = 100$).

information.

Random Problem Instances 2. The RPI2 experiment generated random problem instances that do not match the underlying assumptions of the VIP and OCBA procedures regarding a noninformative distribution of means and variances. In one setting, there are few very good systems ($a = 0$). In another, there are likely to be several good systems ($a = 1$). Procedure $\mathcal{LL}$ and $\mathcal{OCBA}_{\mathrm{LL}}$ again perform almost identically for efficiency and controllability, so only the latter is shown in plots.

Figure 2.20 compares the efficiency and target curves for $\mathcal{OCBA}_{\mathrm{LL}}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$ and $\mathcal{KN}{+}{+}$. Procedures $\mathcal{OCBA}_{\mathrm{LL}}$ and $\mathcal{LL}$ are somewhat less efficient than $\mathcal{KN}{+}{+}$ if there are several good systems ($a = 1$, left panel). The difference is smaller for larger δ^*. Procedures $\mathcal{OCBA}_{\mathrm{LL}}$ and $\mathcal{LL}$ meet the target, and $\mathcal{KN}{+}{+}$ significantly overdelivers $\mathrm{PBS}_{\mathrm{iz},0.2}$. The RPI2 configuration with few good systems ($a = 0$) is quite similar to RPI1 with respect to the long tail distribution of the good systems, so it is not surprising that results are very similar. But even for many good systems ($a = 1$), most of the results from RPI1 carry over (not shown).

On the whole, $\mathcal{OCBA}_{\mathrm{LL}}$ and $\mathcal{LL}$ perform very well for RPI2 even though the problem instances do not follow the normal-inverted gamma distribution that is implicit in the derivation of those procedures. A small degradation in efficiency relative to $\mathcal{KN}{+}{+}$ may be expected if there are multiple very good systems, but controllability remains with $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$. Procedures 0-1,

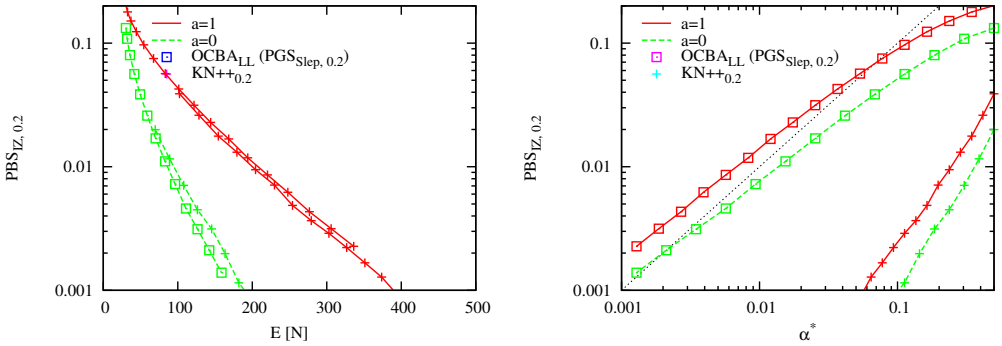

Figure 2.20: Efficiency and target for RPI2 ($k = 5$, $\eta = 1$, $\alpha = 100$). $\mathcal{LL}$ and $\mathcal{OCBA}_{\mathrm{LL}}$ perform almost identically.

0-1_1 and $\mathcal{LL}_1$ are less effective.

Additional Supporting Graphs

The previous section presented a summary of the general conclusions from the study. This section contains a subset of additional results that explore the ideas further. It is not practical to display all results from the experiments, as we tested over 150 problem configurations defined by combinations of k, {SC, MDM, RPI1, RPI2}, and configuration parameters (δ, ρ for SC, MDM; η, α, β for RPI1, RPI2). Together with combinations of n_0, sampling allocations, and stopping rule parameters, over 10^4 different combinations were run. We developed a graphical visualization tool to allow an easy navigation through the results. That tool was used to generate most of the figures in this chapter.

SC, $k > 2$

Figure 2.21 (left panel) illustrates the observation that the advantage of adaptive Bayesian procedures, relative to Equal, increases with k. The qualitative nature of the graph does not change as δ and ρ are individually varied from the values used for the plot. For all other procedures with the new stopping rules, there is a tendency to overdeliver as k increases, but $\mathcal{OCBA}_{\mathrm{LL}}$ is more sensitive than Equal (right panel). The tendency to overdeliver with increasing k might be attributed to the slack introduced by Slepian's and Bonferroni's inequalities.

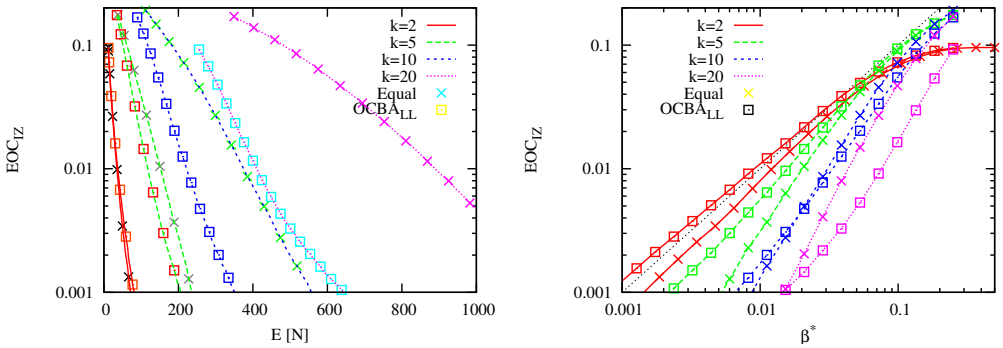

Figure 2.21: Influence of the number of systems k on efficiency (right panel) and target (left panel). Equal and $\mathcal{OCBA}$ allocation are used in combinaton with EOC_{Bonf} stopping rule (SC, $\delta = 0.5$, $\rho = 1$).

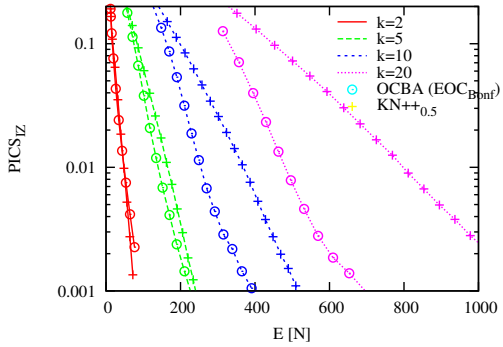

Figure 2.22: Comparison of $\mathcal{KN}{+}{+}$ and $\mathcal{OCBA}(\text{EOC}_{\text{Bonf}})$ efficiency depending on the number of systems k (SC, $\delta = 0.5$, $\rho = 1$).

Similar to Equal, $\mathcal{KN}{+}{+}$ also loses efficiency relative to $\mathcal{OCBA}(\text{EOC}_{\text{Bonf}})$ as the number of systems k increases (Figure 2.22).

Figure 2.23 shows the importance of a sufficiently large n_0 also for $k > 2$. The right panel indicates that increasing n_0 increases the tendency to overdeliver. While $\mathcal{LL}(\text{EOC}_{\text{Bonf}})$ is slightly closer to the target than $\mathcal{OCBA}_{\text{LL}}(\text{EOC}_{\text{Bonf}})$, it is slightly less efficient. $\mathcal{KN}{+}{+}$ is quite insensitive to n_0 (not shown).

A larger variance for the best system relative to the other systems (larger ρ) makes correct selections easier (Figure 2.24, left panel). Different allocation functions respond differently to ρ, as illustrated for $\mathcal{OCBA}$ and $\mathcal{OCBA}_{\text{LL}}$, representatives for the 0-1- and EOC-based allocations. With $\rho = 4$, the EOC-based allocation is clearly superior, but the advantage diminishes for low PICS as ρ decreases, and the PCS-based allocation partially outperforms

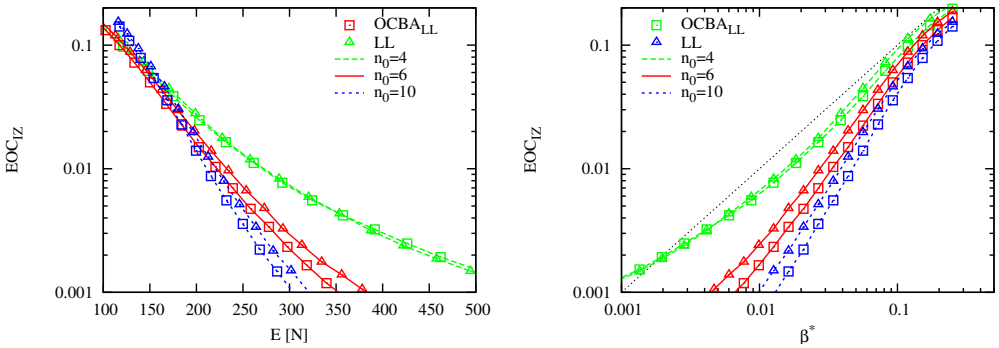

Figure 2.23: Effect of various n_0, with EOC_{Bonf} stopping rule on *efficiency* (left) and *target* (right) (SC, $k = 10$, $\delta = 0.5$, $\rho = 1$).

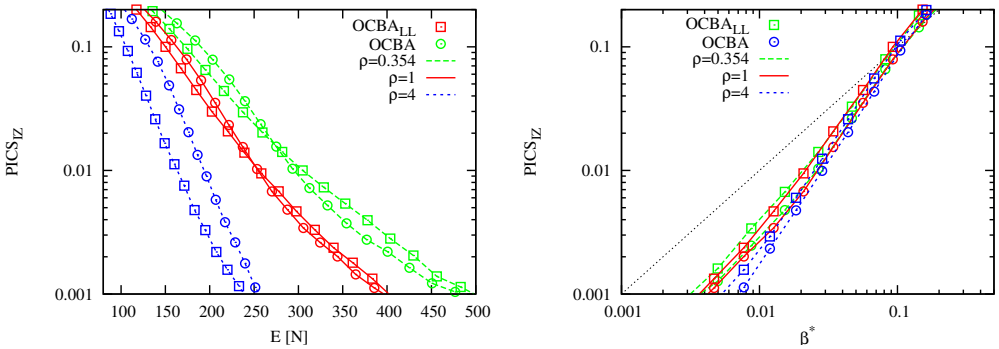

Figure 2.24: Influence of variance ratio ρ on $\mathcal{OCBA}$ and $\mathcal{OCBA}_{LL}$ (SC, $k = 10$, $\delta = 0.5$, EOC_{Bonf}).

the EOC-based allocation. Larger ρ overdeliver slightly more on the target plots, and $\mathcal{OCBA}$ overdelivers slightly more than $\mathcal{OCBA}_{\text{LL}}$ (right panel).

MDM, $k > 2$

Figure 2.25 shows $\text{PBS}_{\text{iZ},\delta^*}$ efficiency and target performance for different δ^*. The parameter δ^* for $\mathcal{KN}$++ and $\text{PGS}_{\text{Slep},\delta^*}$ stopping rule have thereby been set to the indifference zone used for measuring performance. For MDM with $\delta = 0.5$, an indifference zone of $\delta^* = 0.2$ or $\delta^* = 0.4$ are equivalent for $\text{PBS}_{\text{iZ},\delta^*}$ efficiency, as only the best system is considered to be a good selection. For $\delta^* = 0.6$, the two best systems are considered good. On the efficiency plot (left panel), it can be seen that the problem with $\text{PBS}_{\text{iZ},0.6}$ is significantly easier. The efficiency curves for $\delta^* = 0.2$ and 0.4 are very similar

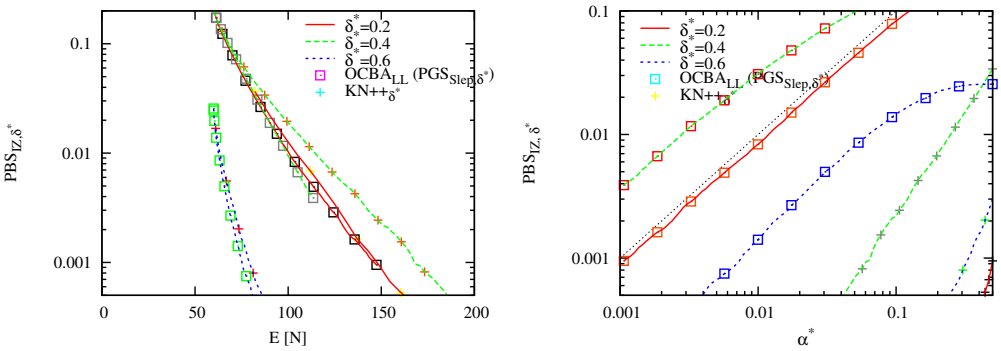

Figure 2.25: $\text{PBS}_{\text{iz},\delta^*}$ efficiency and target for $\mathcal{OCBA}_{\text{LL}}(\text{PGS}_{\text{Slep},\delta^*})$ and $\mathcal{KN}$++ for different settings of δ^* (MDM, $k = 10$, $\delta = 0.5$, $\rho = 1$). In this configuration $\text{PBS}_{\text{iz},0.2} = \text{PBS}_{\text{iz},0.4} = \text{PCS}_{\text{iz}}$.

for $\mathcal{OCBA}_{\text{LL}}$, while $\mathcal{KN}$++ loses efficiency for $\delta^* = 0.4$. On the target plot (right panel), the target performance of both procedures is affected by δ^*. Overall, the $\text{PGS}_{\text{Slep},\delta^*}$ stopping rule is closer to the target than $\mathcal{KN}$++, which very much overdelivers in each case (the curve for $\delta^* = 0.2$ is almost outside the plot).

RPI1

One question is whether the δ^* of the selection procedure can be selected in a clever way to achieve a given desired performance for the probability of a good selection. Figure 2.26 illustrates the influence of δ^* on the $\text{PGS}_{\text{iz},0.2}$ efficiency of $\mathcal{KN}$++ and $\mathcal{OCBA}_{\text{LL}}(\text{PGS}_{\text{Slep},\delta^*})$. For RPI, setting the procedure's parameter δ^* to the δ^* as specified in the efficiency goal (0.2) yields reasonable, though not optimal, efficiency for both procedures. The target performance is good for $\mathcal{OCBA}_{\text{LL}}(\text{PGS}_{\text{Slep},\delta^*})$, while $\mathcal{KN}$++ significantly overdelivers (see Figure 2.17, right panel).

Figure 2.28 extends this observation by looking at different δ^*, but always keeping the δ^* parameter of the procedure equal to the δ^* used to measure efficiency or target performance. The left panel shows that $\mathcal{OCBA}_{\text{LL}}(\text{PGS}_{\text{Slep},\delta^*})$ is more efficient for $\text{PBS}_{\text{iz},\delta^*}$ than $\mathcal{KN}$++ for all settings of δ^*, while the right panel shows that $\mathcal{OCBA}_{\text{LL}}(\text{PGS}_{\text{Slep},\delta^*})$ follows the target better than $\mathcal{KN}$++. The combined effect of lower efficiency and overdelivery is visualized

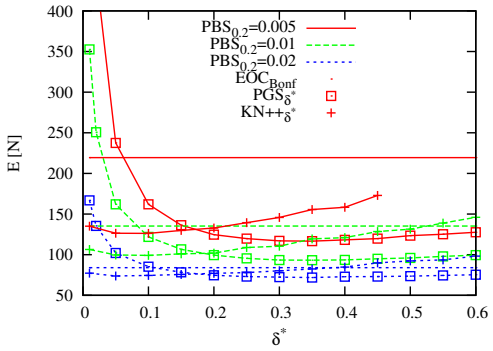

Figure 2.26: Influence of δ^* on the required number of samples to obtain a desired $\mathrm{PBS}_{\mathrm{iz},0.2} = \alpha^*$ for $\mathcal{KN}{+}{+}$ and $\mathcal{OCBA}_{\mathrm{LL}}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$. Horizontal lines show the number of samples required by $\mathcal{OCBA}_{\mathrm{LL}}(\mathrm{EOC}_{\mathrm{Bonf}})$ as reference (RPI1, $k = 5$, $\eta = 1$, $\alpha = 100$).

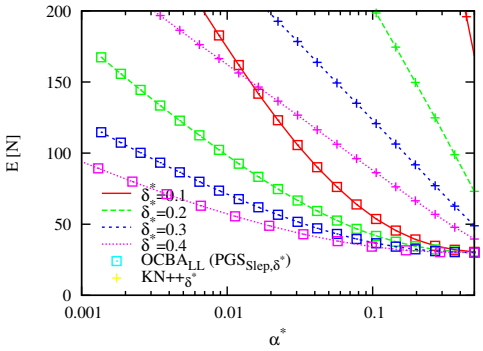

Figure 2.27: Number of samples used by $\mathcal{OCBA}_{\mathrm{LL}}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$ and $\mathcal{KN}{+}{+}$ depend on the setting of parameter α^* (RPI1, $k = 5$, $\eta = 1$, $\alpha = 100$).

in Figure 2.27, which shows the number of samples taken depending on the parameter α^*, again for three different δ^* configurations. For example, if $\delta^* = 0.4$ and a desired accuracy of $\alpha^* = 0.01$ is chosen, $\mathcal{OCBA}_{\mathrm{LL}}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$ uses a total of 57 samples, while $\mathcal{KN}{+}{+}$ uses 161. Both methods deliver at least the desired target accuracy in this example, although for $\mathcal{OCBA}_{\mathrm{LL}}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$, this is not true for all settings of α^* and δ^* (see Figure 2.28 right panel).

Some other observations from SC and MDM also carry over to RPI: The Bayesian procedures are generally quite sensitive to n_0, while $\mathcal{KN}{+}{+}$ is not (Figure 2.29), and $\mathcal{KN}{+}{+}$ becomes less efficient relative to the Bayesian procedures as the number of systems k increases (Figure 2.30).

Figure 2.31 shows that the benefit of including prior information in the VIP and OCBA procedures is more or less independent of α and η for the values tested.

RPI2

In Section 2.5 the paragraph on "'Random Problem Instances 2"' states that most observations made for RPI1 carry over to RPI2 even in the case of

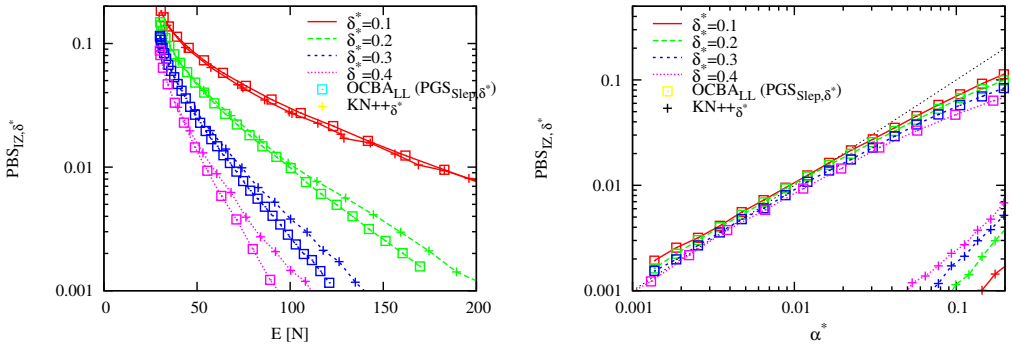

Figure 2.28: $\mathrm{PBS}_{\mathrm{iz},\delta^*}$ efficiency and target for $\mathcal{OCBA}_{\mathrm{LL}}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$ and $\mathcal{KN}{+}{+}$ for different settings of δ^* (RPI1, $k = 5$, $\eta = 1$, $\alpha = 100$).

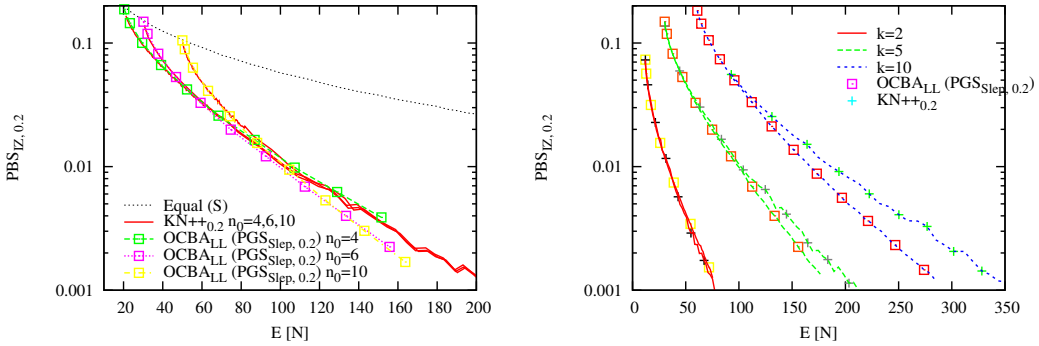

Figure 2.29: Influence of n_0 on the efficiency of $\mathcal{OCBA}_{\mathrm{LL}}$ and $\mathcal{KN}{+}{+}$ (RPI1, $k = 5$, $\eta = 1$, $\alpha = 100$).

Figure 2.30: Efficiency depending on the number of systems k (RPI1, $\eta = 1$, $\alpha = 100$).

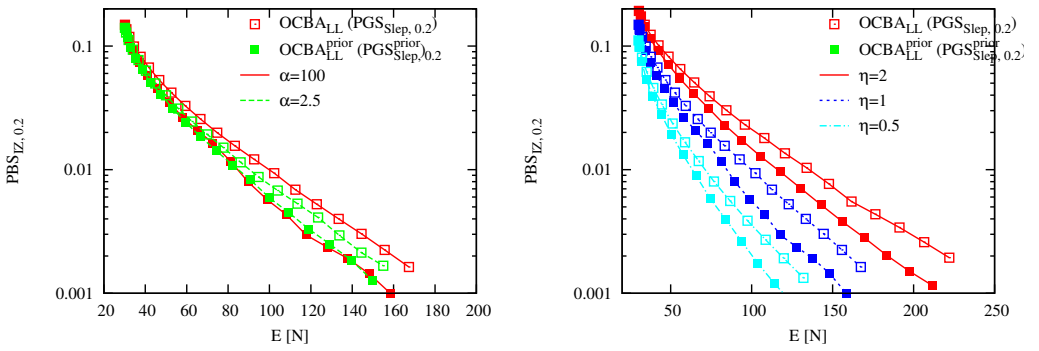

Figure 2.31: Benefit of providing prior information depending on problem configuration parameters α, η (RPI1, $k = 5$, $\alpha = 100$, unless specified otherwise).

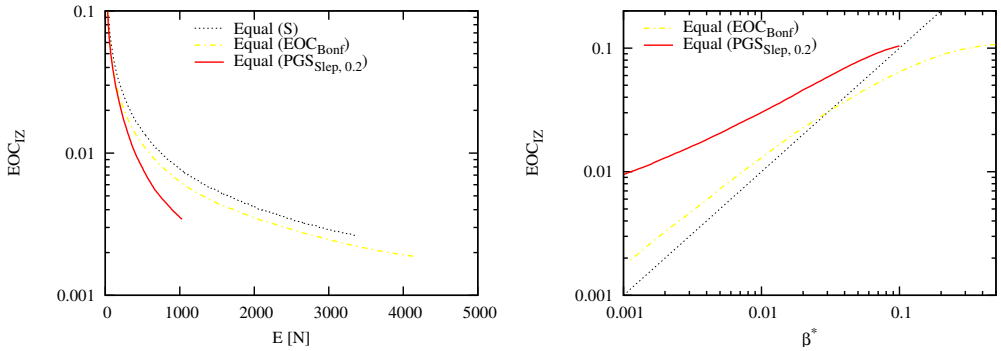

Figure 2.32: Efficiency (left) and target (right) for Equal allocation and different stopping rules w.r.t. $\mathrm{EOC}_{\mathrm{iz}}$ (RPI2, $k = 5$, $a = 1$, $\eta = 1$, $\alpha = 100$). For $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$ stopping, β^* is approximated by $\alpha^*\delta^*$

$a = 1$, i.e. many good systems. Some evidence for this claim is given below.

Figure 2.33 shows that the main conclusions about allocation rules also hold for RPI2. The $\mathcal{LL}$, $\mathcal{OCBA}_{\mathrm{LL}}$ and $\mathcal{OCBA}_{\delta^*}$ allocations are more or less equally efficient. Procedures 0-1 and Equal are clearly less efficient.

The relative ordering of the stopping rules with respect to $\mathrm{EOC}_{\mathrm{iz}}$ efficiency in combination with the Equal allocation remains the same: $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$ is more efficient than $\mathrm{EOC}_{\mathrm{Bonf}}$ which is more efficient than $\mathcal{S}$ stopping rule (Figure 2.32). Figure 2.34 compares target plots for RPI with negative exponential (RPI2, a=1), Gaussian (RPI1) and positive exponential (RPI2, a=0) distribution of the means, in the order of decreasing number of good systems. If the sampling distribution for the means matches the type of prior distribution used for the Bayesian procedures, the target is matched closely. Modifying the distribution towards more good systems (negative exponential) or fewer good systems (positive exponential) leads to over- and underdelivery, respectively.

Implementation Issues Figure 2.35 compares $\mathcal{OCBA}_{\delta^*}$ and $\mathcal{OCBA}_{\mathrm{max},\delta^*}$ with both the $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$ and $\mathrm{PCS}_{\mathrm{Slep},\delta^*}$ stopping rules. The result is typical, namely that $\mathcal{OCBA}_{\delta^*}$ is the better allocation and $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$ is the better stopping rule.

We now turn to two implementation issues. [Chen, Yücesan, Dai, and Chen 2005] wrote that the efficiency of $\mathcal{OCBA}(\mathcal{S})$ was not significantly differ-

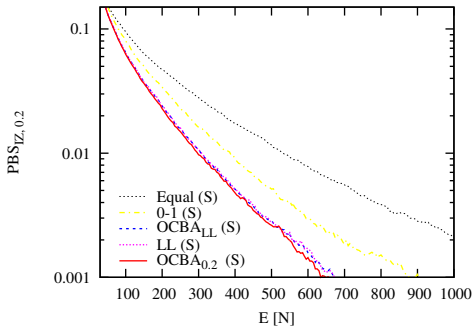

Figure 2.33: $\text{PBS}_{\text{iZ},\delta^*}$ efficiency of allocations with $\mathcal{S}$ stopping rule (RPI2, $k = 5$, $a = 1$, $\eta = 1$, $\alpha = 100$).

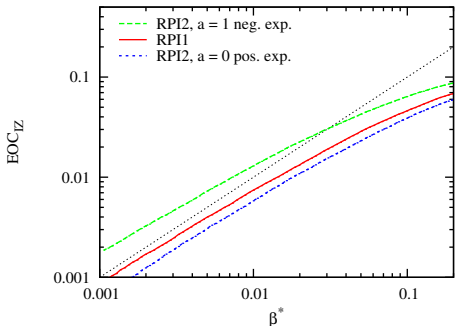

Figure 2.34: Influence of the sampling distributions of configurations on target (RPI1, RPI2, $k = 5, \eta = 1, \alpha = 100$, Equal($\text{EOC}_{\text{Bonf}}$)).

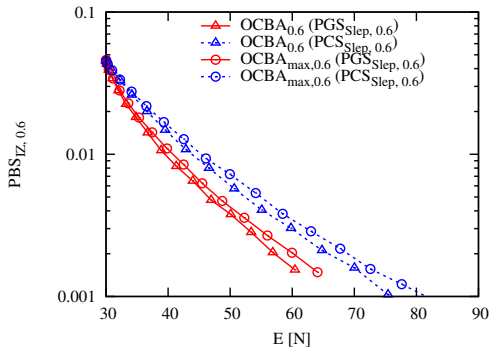

Figure 2.35: Different ways to use δ^* (RPI, $k = 5$, $\eta = 1$, $\alpha = 100$).

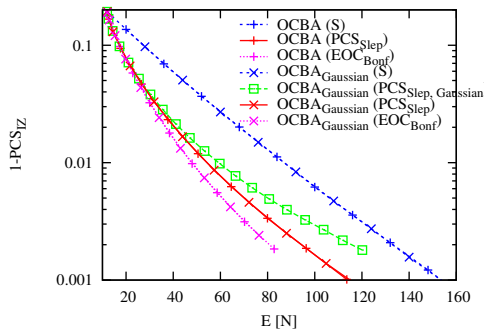

Figure 2.36: Allocation with Normal approximation is as efficient as t, but Normal for stopping rule is less efficient (SC, $k = 2$, $\delta = 0.5$, $\rho = 1$).

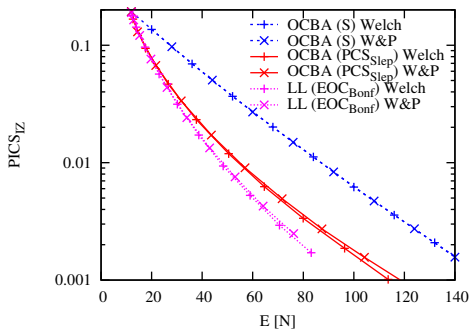

Figure 2.37: Wilson and Pritsker's (W&P) degree of freedom correction was not more efficient than Welch's (SC, $k = 2$, $\delta = 0.5$, $\rho = 1$).

ent whether a t or a normal distribution is used for EAPCS_i (by substituting in the sample variance for the unknown actual variance into a normal distribution version of EAPCS_i), but did not publish results. Figure 2.36 confirms those claims and generalizes to other stopping rules. A normal distribution in the allocation is denoted $\mathcal{OCBA}_{\text{Gaussian}}$. On the other hand, using a normal distribution for the stopping rule ($\text{PCS}_{\text{Slep,Gaussian}}$) *does* degrade performance. The probable cause is that absolute values are important for stopping, but for allocation, relative values for different systems are compared.

A refined estimator of the degrees of freedom that gave good confidence interval coverage for queueing experiments with small numbers of observations ([Wilson and Pritsker 1984]) didn't improve upon Welch's approximation for the SC in Figure 2.37. The associated target plot gave a small (statistically significant) decrease in PCS_{iz} for W&P relative to Welch.

2.6 Discussion

Choices regarding IZ guarantees, probability of correct selection goals, or expected opportunity cost goals make a practical difference in the performance of selection procedures. The new experimental setup (random problem instances as well as stylized configurations) and measurement displays (efficiency and target curves, as opposed to tabular results of PCS_{iz} and the

number of replications), proved useful for assessing the volumes of data, and for identifying strengths and weaknesses of procedures.

For a fixed budget constraint on the number of samples, e.g. due to project deadlines, Procedures $\mathcal{LL}(\mathcal{S})$, $\mathcal{OCBA}_{\mathrm{LL}}(\mathcal{S})$ and $\mathcal{OCBA}_{\delta^*}(\mathcal{S})$ were most efficient.

In the absence of a fixed budget constraint, the results depend on the problem class. For SC and MDM configuration, $\mathcal{LL}$ and $\mathcal{OCBA}_{\mathrm{LL}}$ together with the $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rule were generally the most efficient. $\mathcal{KN}$++ and Wald were also very efficient when $k = 2$, with similar variances and low PICS values. Procedures based on an indifference zone ($\mathcal{KN}$++, Wald and $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$ stopping rule) were very sensitive to the parameter δ^*. While $\mathcal{KN}$++ and Wald followed the target very well for the SC when $\delta^* = \delta$, the difference δ in the best and second best is usually unknown, so this doesn't help in practice. No procedure was particularly controllable for the SC and MDM configurations.

An arbitrary configuration encountered in practice is not likely to be stylized like SC or MDM configurations. When configurations are randomized (RPI1 and RPI2), the $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$ and $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rules can be considered to be reasonably controllable for a desired $\mathrm{PGS}_{\mathrm{iz},\delta^*}$ and $\mathrm{EOC}_{\mathrm{iz}}$, respectively. For RPI, an indifference zone seems important for efficiency. Procedures $\mathcal{LL}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$, $\mathcal{OCBA}_{\mathrm{LL}}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$ and $\mathcal{KN}$++ were the most efficient, with $\mathcal{OCBA}_{\delta^*}(\mathrm{PGS}_{\mathrm{Slep},\delta^*})$ strongly competitive for $\mathrm{PGS}_{\mathrm{iz},\delta^*}$ efficiency. Procedure $\mathcal{KN}$++ tends to overdeliver and is not controllable for $\mathrm{PGS}_{\mathrm{iz},\delta^*}$ in RPI experiments. The current implementation of Wald's procedure cannot be recommended for more than 2 systems.

Strengths of $\mathcal{KN}$++ include a low sensitivity to the number of first stage samples (n_0), and its natural ability to account for correlated output. It is the only method tested here with a proven asymptotic guarantee for $\mathrm{PCS}_{\mathrm{iz}} \geq 1 - \alpha^*$, if that is desired rather than $\mathrm{PCS}_{\mathrm{iz}} = 1 - \alpha^*$. The potential cost of that preference is strong conservativeness and lack of controllability.

We suggest combining the Bayesian allocation procedures with an adaptive stopping rule to substantially improve efficiency. Independent of the stopping rule, the loss-based allocations $\mathcal{LL}$ and $\mathcal{OCBA}_{\mathrm{LL}}$ are among the

most efficient allocations. The most efficient and controllable stopping rule depends on the desired goal (EOC_{Bonf} or $\text{PGS}_{\text{Slep},\delta^*}$). The strong efficiency is relatively robust w.r.t. different configurations, and controllability is relatively robust for RPI. These procedures also allow for the incorporation of prior information about problem instances when that is available. Weak points of those procedures are a dependency on the initial number of samples for each system, n_0, and the potential for a small degradation in performance if there are many systems that are close to the best.

Experiments showed that the Bonferroni and Welch approximations used by $\mathcal{LL}$ are a less significant source of potential suboptimality than the greedy one-step lookahead EVI procedures derived here ($\mathcal{LL}_1$, 0-1_1). Those new procedures are not recommended for general use, nor are Equal and 0-1.

There are several limits to our study. We did not test the effect of autocorrelation from steady-state simulations, but do not see why batching would affect one procedure differently than another. We did not test common random numbers (CRN), a technique that can sharpen contrasts between systems. A strength of $\mathcal{KN}$++ is that it can directly account for CRN. CRN can be accounted for in the OCBA [Fu, Hu, Chen, and Xiong 2005] and VIP [Chick and Inoue 2001a] frameworks, but more work remains. This thesis assumed that the number of samples determines efficiency, but the mean CPU time of different systems may differ in general. The derivation of $\mathcal{LL}$ accounts for different CPU times, and the OCBA is amenable to different sampling costs too. Two-stage procedures are preferred in some contexts. [Inoue, Chick, and Chen 1999] suggests that $\mathcal{LL}$ runs very effectively in two stages. The EOC_{Bonf} stopping rule can be adapted to two stages by finding a second-stage sampling budget that achieves a desired predicted EOC_{Bonf}.

2.7 Future Work

2.7.1 Removing approximations in OCBA

The OCBA procedures are based on the assumption that the mean and variance estimates will not change for a few additional samples. With the

conjugate and noninformative posterior predictive distributions for a single observable X_i' given in [Bernardo and Smith 1994] the posterior distribution of APCS_i for τ_i additional samples can be estimated.

$$p(X_i'|x_{ij} \text{ with } j = 1, \ldots, n_i) = \texttt{St}\left(\bar{x}_i, \frac{n_i - 1}{(n_i + 1)s_i^2}, n_i - 1\right), \tag{2.28}$$

where the mean and variance are $\mathbf{E}[X_i'] = \bar{x}_i$, $\mathbf{Var}[X_i'] = \frac{n_i+1}{n_i-3}s_i^2$. The updated statistics including τ_i new, yet unknown values X_i' are $n_i' = n_i + \tau_i$ independent of X_i', $\bar{x}_i' = \frac{1}{n_i+\tau_i}(n_i\bar{x}_i + \sum^{\tau_i} X_i')$ and $s_i^{2\prime} = \frac{1}{n_i+\tau_i-1}[(n_i - 1)s_i^2 + n_i\bar{x}_i^2 + \sum^{\tau_i} X_i'^2 - \frac{1}{n_i+\tau_i}(n_i\bar{x}_i + \sum^{\tau_i} X_i')^2]$. Their expected values with respect to the distribution of X_i' are $\mathbf{E}_{X_i'}[\bar{x}_i'] = \bar{x}_i$, $\mathbf{E}_{X_i'}[s_i^{2\prime}] = \frac{n_i-1}{n_i+\tau_i-1}s_i^2$.

APCS_i is a function of X_i'. The idea of OCBA is to evaluate the system that has the best PCS_{Slep} after drawing X_i', so we choose the system that maximizes $\mathbf{E}_{X_i'}[\text{APCS}_i]$. OCBA approximates $\mathbf{E}_{X_i'}[\text{APCS}_i] = \text{APCS}_i(\mathbf{E}[X_i'])$. A better approximation should be given by the second taylor polynomial (see e.g. [Rinne 1997])

$$\mathbf{E}_{X_i'}[\text{APCS}_i] \approx \text{APCS}_i(\mathbf{E}[X_i']) + \mathbf{Var}[X_i']\frac{\partial^2}{\partial X_i'^2}\text{APCS}_i(\mathbf{E}[X_i'])$$

Removing this approximation in OCBA should make the allocation more stable in the early phase of the procedure, where estimation of the means and especially the variances are still unstable.

2.7.2 Improving Stopping Rules for Bayesian Procedures

The stopping criteria $\text{PGS}_{\text{Slep},\delta^*}$ and EOC_{Bonf} are realizations of Brownian motion processes with an exponentially decreasing drift.

Currently the procedures stop, if the observed value falls below a given bound. Two problems may arise: 1) The rule is too conservative, i.e. the current decision is correct, but the rule does not indicate to stop. 2) The rule is too optimistic, i.e. the rule stops although the decision is wrong.

The Bayesian procedures with adaptive stopping rules have a curvature

for low values of $\mathrm{PBS}_{\mathrm{iz},\delta^*}$ and a strong dependence on n_0 indicating that at least problem 2 is relevant. A part of the runs stop too early and make errors that cannot be corrected anymore. One improvement would be to add an uncertainty estimator in the stopping rule, so that the procedure stops, if it is certain that the $\mathrm{PBS}_{\mathrm{iz},\delta^*}$ is below a given bound.

2.7.3 Nonnormal Distributions

The procedures presented up to now are based on the order precedence relation $X \prec Y :<=> E(X) < E(Y)$. They can be applied, if the expected values exist. From a practionist's point of view, these procedures might run into troubles, if the variances are very large or the distributions are strongly skewed. If batching several replications does not lead to sufficiently Gaussian distributed replications, there exist two statistical tests, which help in the case of strongly nonnormal distributions: The ranked t-test and Wilcoxon's signed t-test. They are based on the precedence relation $X \prec Y :<=> P(X < Y) > P(X > Y)$. The random variable which delivers a lower value more often is preferred.

For symmetrical distributions this precedence relation is equivalent to the expected value precedence relation. A decision maker will in general base his decisions on the expected value[3]. But if the distribution of the simulation outputs do not have a Gaussian-like distribution at all, one might be forced to change to the order based precedence relation.

2.7.4 Non-Independent Simulations

Many simulations used in practice are steady-state simulations, where the outputs are not independent, but dependent, i.e. the outputs of a simulation are influenced by previous outputs.

Another source of non-independent simulations arises with the use of variance reduction techniques like antithetic random variables, Latin Hypercube Sampling or Quasi Random Numbers.

[3]risk-averse decision makers will base their decision on the expected utility

Kim and Nelson suggest to combine several outputs to a batch, and treat the batched values as approximately independent. Ignoring the existence of non-independent simulations leads to an underestimation of the variance and therefore to too early stopping in statistical selection.

The use of common random numbers (CRN) between systems introduces another desired cause of dependence. In general the procedures presented work in the presence of CRN, but will be too conservative. The efficiency might be improved by estimating the correlation and calculating $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$ with the formulas given in [Genz and Bretz 2002]. $\mathcal{KN}$++ already allows for CRN. The incorporation of variance reduction techniques is one promising area for further research.

Chapter 3

Integrating Statistical Ranking into the Evolutionary Framework

Evolutionary Algorithms (EA) are iterative and probabilistic search algorithms that make use of Darwinian evolution to find good solutions to difficult optimization problems. In each iteration ("generation") from a set of solutions ("population") some individuals are selected as parents ("selection" or "mating selection"). These parents are recombined by crossover and mutated to produce new offspring ("reproduction"). The offspring replaces some or all individuals from the old population to form a new population ("replacement" or "environmental selection"). Better or "fitter" solutions have either a higher probability of being selected as parents, or of surviving to the new population, or both. In the search process, good characteristics (building blocks) emerge in the population. The course of an EA with it's three phases (selection, reproduction, and replacement) can be seen in Figure 3.1.

After each generation, some statistics of the new population are calculated, like best, mean and variance of the fitness values in the population or the diversity of the population. The search stops if one of the statistics fulfills a certain criterion, e.g. the diversity or fitness variance is low. Other

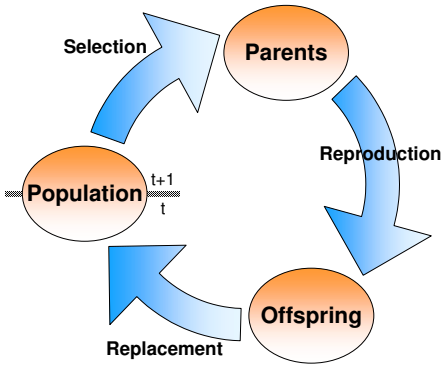

Figure 3.1: Schema of the evolutionary framework

stopping criteria are a maximal number of iterations, evaluations or time overall or without improvement in the best solution's fitness. When using an EA interactively the algorithm can be stopped by the decision maker, as soon as she is satisfied with the result. If the replacement operator ensures that the best solution always survives, then the best solution from the last solution is returned as "the solution". If not, the best solution found so far is saved externally and returned in the end.

An instance of an EA is determined by the choice of a selection operator, recombination operators, a replacement operator, and eventually by a stopping criterion. In the following, the number of individuals in the population is denoted by p, the number of selected parents by m and the number of offspring by o. In the EA literature normally μ and λ are chosen for the population and offspring size. We changed the notation to prevent confusion with mean and precision from the previous chapter. Some examples of algorithms that fit into this scheme are given below.

Evolution Strategies A (p, m)-Evolution Strategy selects m parents from the population of size p (with replacement) with equal probability, mutates each parent to produce new offspring and then the p best offspring form a new population. A $(p + m)$-Evolution Strategy selects the p best out of the

offspring and the old population.

Generational EA Generational EA have a population of size p. To determine the parents, p tournaments of several individuals are made and the winner of each tournament is selected. Two parents are combined by crossover and mutation to form two new offspring. All p offspring replace the parent generation. Often generational EA are combined with elitism, which means that the replacement is modified such that the offspring replace all individuals from the old population except the best.

Steady-State EA Steady-State EA have a population of size p, too, but only two parents are selected per generation to form only one new offspring by crossover and mutation. The new offspring replaces the worst individual in the old population.

Island EA An Island EA consists of several independent populations. The populations interchange individuals from time to time (migration). In non-migrating generations, the islands behave similarly to independent EA. During migration, an additional selection occurs on each island to select the individuals to migrate to the other populations, where they replace the worst individuals.

Simulated Annealing Even Simulated Annealing and Local Search are covered by this framework: The population is of size one, therefore the selection is trivial. The recombination operator produces an offspring in the neighborhood of the parent and then replaces the parent if it's fitness is higher or in the case of Simulated Annealing replaces the parent even for lower fitness, but with decreasing probability.

The algorithms given above are just some typical configurations. In practice several further variants exist. An overview is given for example in [Michalewicz and Fogel 1999].

3.1 Related Work

Many real-world optimization problems are noisy, i.e. a solution's quality (and thus the fitness function) is a random variable. Examples include all applications where the fitness is determined by a randomized computer simulation, or where fitness is measured physically and prone to measuring error.

Already many years ago, researchers have argued that EA should be relatively robust against noise ([Fitzpatrick and Grefenstette 1988]), and recently a number of publications have appeared which support that claim at least partially ([Miller and Goldberg 1996; Arnold and Beyer 2000a; Arnold and Beyer 2000b; Arnold and Beyer 2003]). A recent survey on this topic is [Jin and Branke 2005].

For most noisy optimization problems, the uncertainty in fitness evaluation can be reduced by sampling an individual's fitness several times and using the average as estimate for the true mean fitness. Sampling n times reduces a random variable's standard deviation by a factor of $\sqrt{n}$, but on the other hand increases the computation time by a factor of n. Thus, there is a generally perceived trade-off: either one can use relatively exact estimates but only evaluate a small number of individuals (because a single estimate requires many evaluations), or one can let the algorithm work with relatively crude fitness estimates, but allow for more evaluations (as each estimate requires less effort).

The application of EA in noisy environments has been the focus of many research papers. There are several papers looking at the trade-off between population size and sample size to estimate an individual's fitness, with sometimes conflicting results. [Fitzpatrick and Grefenstette 1988] conclude that for the genetic algorithm studied, it is better to increase the population size than to increase the sample size. On the other hand, [Beyer 1993] shows that for a $(1, o)$ evolution strategy on a simple sphere, one should increase the sample size rather than o. [Hammel and Bäck 1994] confirm these results and empirically show that it also doesn't help to increase the parent population size p. Finally, [Arnold and Beyer 2000a; Arnold and Beyer 2000b] show analytically on the simple sphere that increasing the parent population size

p is helpful in combination with intermediate multi-recombination. [Miller 1997] and [Miller and Goldberg 1996] have developed some simplified theoretical models which allow to simultaneously optimize the population size and the sample size. A good overview of theoretical work on EA applied to noisy optimization problems can be found in [Beyer 2000] or [Arnold 2002].

All papers mentioned so far assume that the sample size is fixed for all individuals. [Aizawa and Wah 1994] were probably the first to suggest that the sample size could be adapted during the run, and suggested two adaptation schemes: increasing the sample size with the generation number, and using a higher sample size for individuals with higher estimated variance. [Albert and Goldberg 2001] look at a slightly different problem, but also conclude that the sample size should increase over the run. For (p, o) or $(p + o)$ selection, [Stagge 1998] has suggested to base the sample size on an individual's probability to be among the p best (and thus should survive to the next generation). [Hedlund and Mollaghasemi 2001] use an indifference-zone selection procedure to select the best p out of o individuals within a genetic algorithm.

For tournament selection, [Branke and Schmidt 2003; Branke and Schmidt 2004] and [Cantu-Paz 2004] use sequential sampling techniques to reduce the number of samples to the minimum required to discriminate between individuals in a tournament. Adaptive sampling strategies have been examined for situations where the noise strength varies over space ([Di Pietro, While, and Barone 2004]). [Boesel 1999] argues that for linear ranking selection, it is sufficient to group individuals of similar quality into one rank, and a corresponding mechanism is proposed.

To "clean up" after optimization (to identify the best, with high probability, of all visited solutions), [Boesel, Nelson, and Kim 2003b] use a ranking and selection procedure after the EA is finished. Recently, [Buchholz and Thümmler 2005] used a statistical selection technique for selection in a $p + o$ strategy as well as to maintain a pool of promising candidates throughout the run from which at the end the final solution is selected.

None of the above works provide the general framework and tight integration of selection algorithms and EA that is suggested here.

A comprehensive overview and extensive comparison of different selection

procedures was given in Chapter 2, which concludes that $\mathcal{OCBA}$ together with an appropriate stopping rule is among the best performing statistical selection procedures. Our approach for adapting EA on noisy problems is based on $\mathcal{OCBA}$.

3.2 A Statistical Ranking Procedure

In general, the reproduction operators are independent of an individual's fitness[1]. So, by noise, only selection and replacement operators are affected. Generally rank based selection and replacement operators are considered to perform better than fitness proportional ones (see [Whitley 1989]). If we define a correct functioning of the EA as making the same selection and replacement decisions as in the deterministic case, only the ranks of individuals in the union of the set of offspring and the set of individuals in the old population need to be determined (in this context the sets can have duplicates, so they are not sets in the mathematical sense, but families). In the deterministic case, the ordering of all individuals can be given easily after the evaluation of each individual. In the stochastic case, the ordering can only be given with a certain degree of certainty. Note that after reproduction the p individuals from the old population have already been evaluated several times, whereas the o offspring are not yet evaluated. The procedures in Chapter 2 can be used repeatedly to obtain the $p+o$ ranks of each individual one by one, but this might be very expensive in terms of samples. Not all procedures can benefit from preevaluated systems.

In this section we derive a procedure for the complete ranking of a set of individuals. Let the size of this set be k. Analogously to the probability of correct or good *selection* (PCS, PGS) and the expected opportunity cost (EOC), let the probability of correct or good *ranking* (PCR, PGR) and the

[1]Exceptions are memetic algorithms, which add hillclimbing after recombination and mutation and some kinds of evolution strategies, where the mutation itself is biased towards promising areas of the search space and not only selection and/or replacement. For the latter case, often the estimates for the fitness values from the selection phase suffice. The evaluations needed for the hillclimbing in the first case will not be integrated into the mechanism presented.

expected opportunity cost of a ranking (EOCR) for the frequentist perspective (IZ), given the configuration $\boldsymbol{\chi} = (\boldsymbol{\mu}, \boldsymbol{\sigma^2})$ be defined as

$$\begin{aligned}
\mathrm{PCR}_{\mathrm{iz}} &\stackrel{\mathrm{def}}{=} \Pr\left(\bigwedge_{i=1}^{k} \mu_{(i)} = \mu_{[i]} \,|\, \boldsymbol{\chi}\right) \\
\mathrm{PGR}_{\mathrm{iz},\delta^*} &\stackrel{\mathrm{def}}{=} \Pr\left(\bigwedge_{i=1}^{k} \mu_{(i)} \geq \mu_{[i]} - \delta^* \,|\, \boldsymbol{\chi}\right) \\
\mathrm{EOC}_{\mathrm{iz}}^{\mathrm{Rank}} &\stackrel{\mathrm{def}}{=} \sum_{i=1}^{k} \mathbf{E}\left[\mu_{[i]} - \mu_{(i)} \,|\, \boldsymbol{\chi}\right]
\end{aligned} \tag{3.1}$$

and the Bayesian perspective, given the data seen so far $\mathcal{E}$, respectively

$$\begin{aligned}
\mathrm{PCR}_{\mathrm{Bayes}} &\stackrel{\mathrm{def}}{=} \Pr\left(\bigwedge_{i=1}^{k} M_{(i)} \geq M_{<i;k>} \,|\, \mathcal{E}\right) \\
\mathrm{PGR}_{\mathrm{Bayes}} &\stackrel{\mathrm{def}}{=} \Pr\left(\bigwedge_{i=1}^{k} M_{(i)} + \delta^* \geq M_{<i;k>} \,|\, \mathcal{E}\right) \\
\mathrm{EOC}_{\mathrm{Bayes}}^{\mathrm{Rank}} &\stackrel{\mathrm{def}}{=} \sum_{i=1}^{k} \mathbf{E}\left[M_{<i;k>} - M_{(i)} \,|\, \mathcal{E}\right],
\end{aligned} \tag{3.2}$$

where $M_{<i;k>}$ denotes the i-th order statistic[2], (i) the order based on the observed means $\bar{x}_i$ and $[i]$ the order based on true means μ_i. The OCBA procedures can easily be adopted for the ranking problem, given the approximations to the Bayesian measures above and the stochastic dominance of $\max_{j<i} M_{(j)} \succeq M_{<i;k>}$[3],

[2]The i-th order statistic is the distribution of the i-th best value out of the observations of k random variables ([Rinne 1997]).

[3]$\max_{j<i} M_{(j)}$ is just one combination out of all possible combinations for the i-th order statistic, so $\Pr(\max_{j<i} M_{(j)} \leq x) < \Pr(M_{<i;k>} \leq x) \forall x$

$$\begin{aligned}
\mathrm{PGR}_{\mathrm{Bayes}} &\geq \prod_{i=1}^{k}\prod_{j<i} \Pr\left(M_{(i)} + \delta^* > M_{(j)} \,|\, \mathcal{E}\right) && (3.3)\\
&\approx \prod_{i=1}^{k}\prod_{j<i} \Phi_{\nu_{(i)(j)}}(\lambda_{ij}^{1/2}(\delta^* + d_{(i)(j)})) = \mathrm{PGR}_{\mathrm{Slep},\delta^*}\\
\mathrm{EOC}_{\mathrm{Bayes}}^{\mathrm{Rank}} &\leq \sum_{i=1}^{k}\sum_{j<i} \mathbf{E}\left[M_{(j)} - M_{(i)} \,|\, M_{(j)} > M_{(i)}\right] && (3.4)\\
&\approx \sum_{i=1}^{k}\sum_{j<i} \lambda_{ij}^{-1/2}\Psi_{\nu_{(i)(j)}}\left[d_{ij}^*\right] = \mathrm{EOC}_{\mathrm{Bonf}}^{\mathrm{Rank}}.
\end{aligned}$$

We denote the OCBA procedure based on $\mathrm{PGR}_{\mathrm{Slep},\delta^*}$-allocation $\mathcal{OCBA}_{\delta^*}^{\mathrm{Rank}}$ and on $\mathrm{EOC}_{\mathrm{Bonf}}^{\mathrm{Rank}}$-allocation $\mathcal{OCBA}_{\mathrm{LL}}^{\mathrm{Rank}}$. PICR and PBR denote the probabilities for incorrect and bad rankings analogously to the definitions in Chapter 2. $\mathrm{PGR}_{\mathrm{Slep},\delta^*}$ and $\mathrm{EOC}_{\mathrm{Bonf}}^{\mathrm{Rank}}$ can be used as stopping rules, too. Furthermore, OCBA can benefit from preevaluated individuals, which we do have from the old population as noted above.

Figure 3.2 shows that efficiency and target for the ranking variant of OCBA are worse than $\mathcal{OCBA}$ for the selection problem – which was expected, because the problem of finding a correct ranking includes the problem of finding the best solution as a subproblem. Nevertheless, $\mathcal{OCBA}_{\delta^*}^{\mathrm{Rank}}(\mathrm{PGR}_{\mathrm{Slep},\delta^*})$ and $\mathcal{OCBA}_{\mathrm{LL}}^{\mathrm{Rank}}(\mathrm{EOC}_{\mathrm{Bonf}}^{\mathrm{Rank}})$ allow to determine a correct ranking with higher efficiency than equal allocation and allow to adapt the number of samples for a given error probability.

Actually an EA does not need a complete ranking for the correct functioning. It suffices, if only a certain set of comparisons are correct, independently whether all ranks are correct. Denote this set of comparisons by $\mathcal{C}$. Which comparisons are needed depends on the chosen selection and replacement operators of the EA. A comparison is given by the tuple $\langle i, j\rangle$, where $\forall\langle i, j\rangle \in \mathcal{C} : i, j \in \{1, \ldots, k\} \wedge i \neq j$. If the observed rank of individual i is higher/lower than that of individual j and i is truly better/worse than j, then the comparison $\langle i, j\rangle$ is correct.

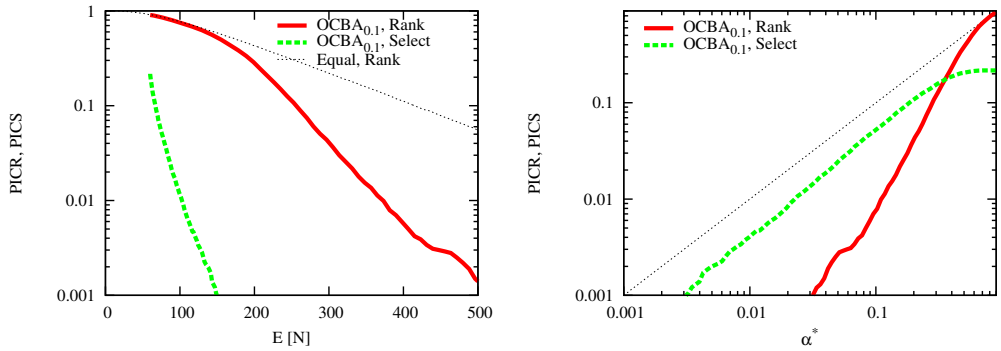

Figure 3.2: Efficiency and target for $\mathcal{OCBA}_{\delta*}^{\text{Rank}}$, $\mathcal{OCBA}$ and Equal allocation for 10 systems, equally spaced ($\delta = 0.5$) and equal variance.

In the following sections, we give the necessary comparisons that need to be added to $\mathcal{C}$ for some popular replacement and selection operators.

3.3 Replacement Operators

The replacement step determines, which p individuals out of the $k = p + o$ individuals from the old population and the offspring make it into the new population. Each individual can be at most once in the new population $\mathcal{P}$[4]. The individuals in the old population are numbered $\mathcal{P}' = \{1, \ldots, p\}$ and the offspring $\mathcal{O} = \{p + 1, \ldots, k\}$. The new population is a true subset $\mathcal{P} \subset \{1, \ldots, k\}$. The order $(.)$ over all individuals is defined by $\bar{x}_{(1)} \leq \ldots \leq \bar{x}_{(k)}$, $(.)_{\mathcal{O}}$ over all offspring and $(.)_{\mathcal{P}}$ over all individuals in the new population. The ranks $(.)^{-1}$ are defined as the inversion of the order ($(r)^{-1} = r$). The orders $(i)_{\mathcal{O}}$ and $(i)_{\mathcal{P}}$ can be determined from the overall order $(.)$ by finding the i-th individual having $(j) \in \mathcal{O}$ or $(j) \in \mathcal{P}$ for $j = 1, \ldots, k$.

3.3.1 Comma- and Plus-Replacement

In (p, o)- and $(p + o)$-replacement simply the p best individuals out of the o offspring or out of the $p+o$ individuals from the old population and offspring are selected. Typically comma- and plus-evolution strategies make use of

[4]Nevertheless we allow for individuals with identical genes in the population

this replacement operator, but steady-state EA use $(p+1)$-replacement and local or tabu search use $(1+1)$-replacement, too.

The set operations for (p, o)-replacement with $o > p \geq 1$ are

$$\begin{aligned} \mathcal{C} &\leftarrow \mathcal{C} \cup \bigcup_{i=o-p+1}^{o} \{\langle (i)_{\mathcal{O}}, (1)_{\mathcal{O}} \rangle, \ldots, \langle (i)_{\mathcal{O}}, (o-p)_{\mathcal{O}} \rangle\} \\ \mathcal{P} &= \{(o-p+1)_{\mathcal{O}}, \ldots, (o)_{\mathcal{O}}\}, \end{aligned}$$

for $(p+o)$-replacement with $p, o \geq 1$

$$\begin{aligned} \mathcal{C} &\leftarrow \mathcal{C} \cup \bigcup_{i=o+1}^{p+o} \{\langle (i), (1) \rangle, \ldots, \langle (i), (o) \rangle\} \\ \mathcal{P} &= \{(o+1), \ldots, (p+o)\}, \end{aligned}$$

and are visualized in a Table 3.1.

Individual i	Individual j				
	(1)	$\ldots$	$(o-p)_{\mathcal{O}}$	$\ldots$	$(p)_{\mathcal{O}}$
$(1)_{\mathcal{O}}$					
$\vdots$					
$(o-p+1)_{\mathcal{O}}$	$\times$	$\cdots$	$\times$		
$\vdots$	$\vdots$	$\ddots$	$\vdots$		
$(p)_{\mathcal{O}}$	$\times$	$\cdots$	$\times$		

Individual i	Individual j				
	(1)	$\ldots$	(o)	$\ldots$	$(o+p)$
(1)					
$\vdots$					
$(o+1)$	$\times$	$\cdots$	$\times$		
$\vdots$	$\vdots$	$\ddots$	$\vdots$		
$(o+p)$	$\times$	$\cdots$	$\times$		

Table 3.1: Comparisons for comma- (left) and plus-replacement (right). A "$\times$" means that the comparison between individuals i and j is needed, $\langle i, j \rangle \in \mathcal{C}$. Note that for comma-replacement the comparisons are only among the offspring, while for plus-replacement for the old population and the offspring.

Plus-replacement ensures that the best solution found so far always survives to the next iteration, while comma-replacement will replace it's best solution every iteration.

3.3.2 Complete Replacement and Elitism

Generational-EA replace the old population completely by the p offspring produced, whereas generational-EA with elitism only produce $p-1$ offspring

and replace all but the best in the old population. Without elitism, the best solution found so far might get lost and thus the generational-EA is mostly combined with elitism.

For complete replacement no comparisons are needed, therefore $\mathcal{C}$ is not modified at all and the new population is the set of the p offspring

$$\mathcal{P} \;=\; \{p+1,\ldots,2p\}\,.$$

Generational-EA with elitism add comparisons between the best and all others in the new population. We determine the elite individual $(p)_{\mathcal{P}}$ for the next iteration in the current iteration, so we can benefit from other comparisons needed in the current iteration due to selection. The number of offspring generated is $p-1$ and the elitist individual is $(p)_{\mathcal{P}'}$.

$$\begin{aligned}\mathcal{P} &= \{(p)_{\mathcal{P}'}, p+1,\ldots,2p-1\}\\ \mathcal{C} &\leftarrow \mathcal{C} \cup \bigcup_{i=1}^{p-1} \{\langle (p)_{\mathcal{P}}, (i)_{\mathcal{P}}\rangle\}\end{aligned}$$

3.4 Selection Operators

The selection step determines which m individuals from the new population $\mathcal{P}$ are selected as parents. The set of parents is denoted $\mathcal{M}$ and, as each individual may be selected several times, can have duplicates.

For the selection operator of an EA several implementations exist. The rank-based determine a probability p_r for the selection of an individual with rank r. This rank is defined with respect to the new population with size p. As not all replacement operators select the top p individuals out of the o offspring and p individuals in the old population, the rank used for determination of individual i's selection probability does not necessarily equal $(i)^{-1}$. Instead we denote the new ranks related to the new population by $(i)^{-1}_{\mathcal{P}}$. Only comparing individuals i and j is not affected by the new ranks, as $(i)^{-1}_{\mathcal{P}} > (j)^{-1}_{\mathcal{P}}$ follows directly from $(i)^{-1} > (j)^{-1}$. So, only selection operators that explicitly rely on the rank are affected.

The expected number of times an individual i is selected can be calculated by multiplying $p_{(i)_{\mathcal{P}}^{-1}}$ with the number of overall selections per generation m. In general, selections are independent of each other, so the number of selections follow a multinomial distribution.

3.4.1 Random Selection

Each individual has the same probability of being selected. This operator is normally used in Evolution Strategies. The optimization pressure is only introduced by the replacement operator. The individuals do not have to be evaluated for random selection, but can be just randomly selected from the population $\mathcal{P}$ independently of the fitness.

$$p_i = \frac{1}{p}$$

Random selection therefore does not add any comparisons to $\mathcal{C}$, and $\mathcal{M}$ is the result of m random samples from $\mathcal{P}$. Denote a random sample of $\mathcal{P}$ by π_j then

$$\mathcal{M} = \{\pi_1, \ldots, \pi_m\}.$$

To apply stochastic universal sampling, each individual in $\mathcal{P}$ is selected $\lfloor m/p \rfloor$-times and the remaining $m - p\lfloor m/p \rfloor$ parents are determined by choosing randomly from $\mathcal{P}$ without replacement, i.e. without duplicates.

3.4.2 Tournament Selection

Tournament selection (TS, see [Goldberg, Korb, and Deb 1989]) chooses two individuals randomly from the population $\mathcal{P}$ and selects the better of the two as parent. This selection operator is mostly used with generational- and steady-state-EA. For generational-EA m parents are selected by m tournament selections and only two tournaments are performed to select two parents for steady-state-EA.

A variant of tournament selection is t-tournament selection, where the parent is determined as the best out of t instead of two randomly chosen individuals. The t individuals are chosen from $\mathcal{P}$ without replacement, so that

an individual cannot participate in a tournament with itself. The selection probability of individual $i \in \mathcal{P}$ with rank r is

$$p_r = \begin{cases} r \geq t : & \frac{(r-1)!(p-t)!}{(r-t)!p!}t \\ r < t : & 0 \end{cases} .$$

Denote the chosen t tournament participants of tournament j by $\pi_j^1, \ldots, \pi_j^t$ and by $\pi_j^{(i)}$ the ordered tournament participants, where $\bar{x}_{\pi_j^{(1)}} < \ldots < \bar{x}_{\pi_j^{(t)}}$. Then the set of comparisons is extended as follows and the selected parents are

$$\begin{aligned} \mathcal{C} &\leftarrow \mathcal{C} \cup \bigcup_{j=1}^{m} \left\{ \left\langle \pi_j^{(t)}, \pi_j^{(1)} \right\rangle, \ldots, \left\langle \pi_j^{(t)}, \pi_j^{(t-1)} \right\rangle \right\} \\ \mathcal{M} &= \{\pi_1^{(t)}, \ldots, \pi_m^{(t)}\} \end{aligned}$$

An example of the comparisons for tournament selection is given in Table 3.2.

Individual i	Individual j									
	$(1)_\mathcal{P}$	$(2)_\mathcal{P}$	$(3)_\mathcal{P}$	$(4)_\mathcal{P}$	$(5)_\mathcal{P}$	$(6)_\mathcal{P}$	$(7)_\mathcal{P}$	$(8)_\mathcal{P}$	$(9)_\mathcal{P}$	$(10)_\mathcal{P}$
$(1)_\mathcal{P}$										
$(2)_\mathcal{P}$										
$(3)_\mathcal{P}$										
$(4)_\mathcal{P}$										
$(5)_\mathcal{P}$										
$(6)_\mathcal{P}$										
$(7)_\mathcal{P}$		×	×		×	×				
$(8)_\mathcal{P}$		×	×	×		×	×			
$(9)_\mathcal{P}$			×				×			
$(10)_\mathcal{P}$	×		×		×	×	×			

Table 3.2: Comparisons for 10 randomly chosen 3-tournaments ($\{1, 6, 10\}$, $\{3, 6, 7\}$, $\{3, 7, 9\}$, $\{2, 7, 8\}$, $\{3, 6, 7\}$, $\{5, 7, 10\}$, $\{2, 5, 7\}$, $\{3, 6, 10\}$, $\{3, 7, 8\}$, $\{4, 6, 8\}$), assuming $(i)_\mathcal{P} = i \in \mathcal{P}$. The set of parents is $\mathcal{M} = \{7, 7, 7, 8, 8, 8, 9, 10, 10, 10\}$.

Local selection is a selection operator where the best individual from a neighborhood of size t is selected. For the statistical ranking procedure, this

is just a special form of t-tournament selection, where the tournament participants are not chosen randomly, but the individuals in the neighborhood.

3.4.3 Stochastic Tournament Selection

Stochastic tournament selection (STS, [Goldberg and Deb 1991]) is a rather simple selection scheme where two individuals are randomly chosen from the population, and then the better is selected with probability $(1-\gamma), \gamma \in [0, 0.5]$. If the p individuals in the population are sorted from rank $r = 1$ (worst) to rank $r = p$ (best), the expected probability p_r to select an individual with rank r can be calculated as

$$p_r = \frac{2}{p}\left(1 - \gamma - (1-2\gamma)\frac{p-r}{p-1}\right).$$

This corresponds to a linearly increasing selection probability depending on an individual's rank r, with the slope of the line being determined by the selection probability $(1-\gamma)$ (cf. Figure 3.3).

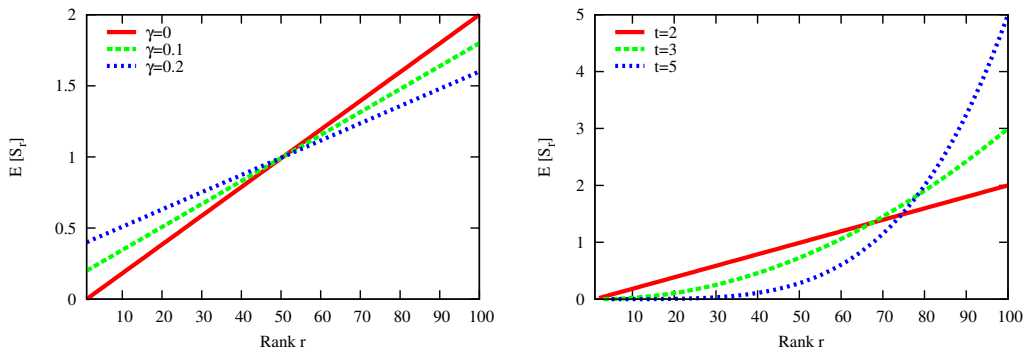

Figure 3.3: Expected number of selections $\mathbf{E}[S_r]$ of an individual with rank r for $p = 100$ individuals and p tournament selections; STS for $\gamma \in \{0, 0.1, 0.2\}$ (left) and TS for $t \in \{2, 3, 5\}$ (right).

In this respect, STS is equivalent to linear ranking selection. But while ranking selection is based on a full ordering of the individuals, STS only requires pairwise comparisons and is easier to implement. For each stochastic tournament i the comparison between two randomly chosen tournament

participants π_i^1 and π_i^2 is needed. The participant with the higher rank, individual $\pi_i^{(2)}$ is selected with probability $1-\gamma$ and $\pi_i^{(1)}$ with probability γ. The same behavior can be achieved by selecting one of the participants randomly with probability 2γ and the better of the two with probability $1-2\gamma$. So for STS the following comparison is required only with probability $1-2\gamma$.

$$\mathcal{C} \leftarrow \mathcal{C} \bigcup_{i=1}^{m} \{(\pi_i^{(2)}, \pi_i^{(1)})\}$$

In Chapter 4 a method is given that increases the accuracy of STS in a stochastic environment by using the noise in the evaluation to replace partially the randomization introduced by selecting the better only with probability $1-\gamma$.

3.4.4 Linear Ranking

Linear ranking selects individual r with linearly increasing probability in r

$$p_r = \frac{1}{p}\left(2 - t + 2(t-1)\frac{r-1}{p-1}\right),$$

with $t \in [1, 2]$.

To calculate the selection probabilities explicitly, a complete ranking is needed, but stochastic tournament selection (see Section 3.4.3) with $\gamma = 0.5(2-t)$ implements the same selection probabilities without the need of a complete ranking.

3.4.5 Exponential Ranking

For exponential ranking the selection pressure is higher than for linear ranking. The selection probabilities are exponentially decreasing in the rank r

$$p_r = \frac{1}{Z} c^{p-r}$$

with $c \in (0, 1)$ and $Z = \frac{1-c^p}{1-c}$.

To exactly match the desired selection probabilities, the complete ranking

needs to be determined. If an approximation to the selection probabilities is acceptable, then exponential ranking can be replaced by t-tournament selection with $t = \left\lfloor p\frac{1-c}{1-c^p} \right\rfloor$, which only needs $t-1$ comparisons per selection.

3.4.6 Further Selection Operators

Truncation selection selects the m parents randomly (with replacement) among the t best individuals from the population. Finding the t best individuals needs the following comparisons

$$\mathcal{C} \leftarrow \mathcal{C} \bigcup_{i=p-t+1}^{p} \{\langle (i)_{\mathcal{P}}, (1)_{\mathcal{P}} \rangle, \ldots, \langle (i)_{\mathcal{P}}, (t)_{\mathcal{P}} \rangle\}\,.$$

Selection operators that need an explicit ordering of all individuals in a population are very expensive, in terms of evaluations and should therefore be avoided. Nevertheless the comparisons are given

$$\mathcal{C} \leftarrow \mathcal{C} \bigcup_{i=1}^{p} \{\langle (i)_{\mathcal{P}}, (1)_{\mathcal{P}} \rangle, \ldots, \langle (i)_{\mathcal{P}}, (i-1)_{\mathcal{P}} \rangle\}$$

3.4.7 Stochastic Universal Sampling

One of the drawbacks of tournament selection is its possibly high sampling error: the *actual* number of times an individual is selected may differ significantly from the *expected* number of times. For example, if one would like to select m parents, and the probability to select an individual r is p_r, then the expected number of times that individual is chosen is $\mathbf{E}[S_r] = m \cdot p_r$, but the actual number s_r can be anywhere between 0 and m for tournament selection (TS).

Such a sampling error is undesired and is often also called genetic drift. The sampling error can be minimized by selecting the parents not independently, but instead using a method called Stochastic Universal Sampling (SUS, [Baker 1987]). If p_r is the probability for selecting individual with rank r, then SUS ensures that the actual number of selections s_r is within $[\lfloor m \cdot p_r \rfloor, \lceil m \cdot p_r \rceil]$.

The effect can be seen in Figure 3.4, which shows the standard deviation of the number of times an individual is selected depending on its rank in a population of 100 individuals. As can be seen, the standard deviation is much lower for SUS than for the standard random TS.

However, the standard deviation for t-TS can be reduced as well by making sure that every individual participates in exactly t tournaments, i.e. the individuals for the tournaments are no longer chosen independently, but are generated by latin rectangles (see [Byers 1991]). Hence this selection scheme is denoted latin TS. Latin STS is implemented by a latin 2-TS, where the first participant is chosen with probability 2γ and the better of the two with probability $1 - 2\gamma$.

Latin TS and STS have exactly the same expected behavior as the standard random schemes, but reduce the selection variance significantly, in particular for the better individuals (see Figure 3.4). For worse individuals, the standard deviation of all three sampling schemes is almost equally. For 2-TS, latin TS and SUS are identically (see dotted line in left figure: the crosses of latin TS are inside the boxes of SUS), but for higher (right figure) or lower selection pressure ($\gamma = 0.1$, left figure) latin TS is not able to reduce the standarddeviation into the deep valleys of SUS except for the highest rank. Nevertheless, the standarddeviation for latin TS is clearly below random TS and almost equal to SUS.

In the above sections, it was shown that the selection probabilities of most rank-based selection schemes can be achieved by tournament selection. The use of latin tournament selection allows to reduce the sampling variance similarly to SUS without the need for further comparisons.

3.5 A Statistical Ranking Procedure for Evolutionary Algorithms

One of OCBA's advantages is its flexibility. It can be easily adapted to not only select the best out of a given set of individuals or determine a complete ranking, but for arbitrary quality measures. To integrate it into an EA,

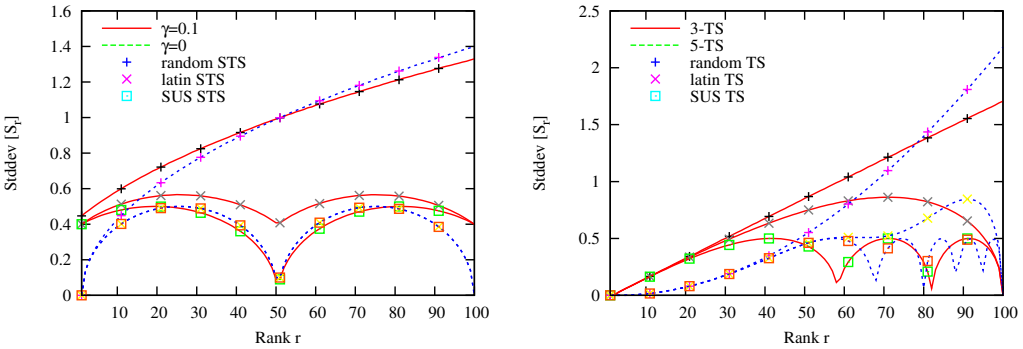

Figure 3.4: Standard deviation of number of selections S_r for different sampling schemes for $p = 100$ individuals and p tournament selections; STS with $\gamma \in \{0, 0.1\}$ (left) and t-TS with $t \in \{3, 5\}$ (right).

we want to make sure that the EA operates "correctly", meaning that the order relations required by the EA's selection and replacement operators have been determined correctly with a sufficiently high confidence. The previous sections explained how to determine the required set of comparisons $\mathcal{C}$. As a quality criterion, we define the probability of good generation ($\mathrm{PGG}_{\mathrm{Bayes}}$) as probability that all pairwise comparisons in $\mathcal{C}$ are correct. The following equation approximates the probability that for all pairs in $\mathcal{C}$, the individual with the higher observed rank also has a higher true mean value. It assumes rank $i >$ rank j for all $\langle i, j\rangle \in \mathcal{C}$. If rank $j >$ rank i, simply replace $\langle i, j\rangle$ by $\langle j, i\rangle$ in $\mathcal{C}$ before calculation.

$$
\begin{aligned}
\mathrm{PGG}_{\mathrm{iz},\delta^*} &\stackrel{\mathrm{def}}{=} \Pr\left(\bigwedge_{\langle i,j\rangle\in\mathcal{C}} \mu_i \geq \mu_j - \delta^* \,|\, \boldsymbol{\chi}\right) \\
\mathrm{EOC}^{\mathrm{Gen}}_{\mathrm{iz}} &\stackrel{\mathrm{def}}{=} \sum_{\langle i,j\rangle\in\mathcal{C}} \mathbf{E}\left[\mu_j - \mu_i \,|\, \mu_j > \mu_i, \boldsymbol{\chi}\right] \qquad (3.5) \\
\mathrm{PGG}_{\mathrm{Bayes}} &\stackrel{\mathrm{def}}{=} \Pr\left(\bigwedge_{\langle i,j\rangle\in\mathcal{C}} M_i \geq M_j - \delta^* \,|\, \mathcal{E}\right) \\
\mathrm{EOC}^{\mathrm{Gen}}_{\mathrm{Bayes}} &\stackrel{\mathrm{def}}{=} \sum_{\langle i,j\rangle\in\mathcal{C}} \mathbf{E}\left[M_j - M_i \,|\, M_j > M_i, \mathcal{E}\right]. \qquad (3.6)
\end{aligned}
$$

Note that the definitions of $\mathrm{EOC}^{\mathrm{Gen}}$ slightly differ from EOC and $\mathrm{EOC}^{\mathrm{Rank}}$. They are based on the sum of the pairwise losses and not the rank wise losses.

As for $\mathcal{OCBA}^{\mathrm{Rank}}_{\delta^*}$ and $\mathcal{OCBA}^{\mathrm{Rank}}_{\mathrm{LL}}$, the Bayesian measures can be approximated by Slepian and Bonferroni bounds and we obtain criteria for the $\mathcal{OCBA}^{\mathrm{EA}}_{\delta^*}$ and $\mathcal{OCBA}^{\mathrm{EA}}_{\mathrm{LL}}$ variants of $\mathcal{OCBA}$ for the ranking needed for EA.

$$\begin{aligned}
\mathrm{PGG}_{\mathrm{Bayes}} &\geq \prod_{\langle i,j\rangle\in\mathcal{C}} \Pr\left(M_i+\delta^* > M_j \,|\, \mathcal{E}\right) \\
&\approx \prod_{\langle i,j\rangle\in\mathcal{C}} \Phi_{\nu_{ij}}(\lambda_{ij}^{1/2}(\delta^*+d_{ij})) = \mathrm{PGG}_{\mathrm{Slep},\delta^*} \\
\mathrm{EOC}^{\mathrm{Gen}}_{\mathrm{Bayes}} &\leq \sum_{\langle i,j\rangle\in\mathcal{C}} \mathbf{E}\left[M_j - M_i \,|\, M_j > M_i\right] \\
&\approx \sum_{\langle i,j\rangle\in\mathcal{C}} \lambda_{ij}^{-1/2}\Psi_{\nu_{ij}}\left[d^*_{ij}\right] = \mathrm{EOC}^{\mathrm{Gen}}_{\mathrm{Bonf}}. \qquad (3.7)
\end{aligned}$$

The procedures are outlined below.

Procedures $\mathcal{OCBA}^{\mathbf{EA}}_{\delta^*}(\alpha^\star)$ and $\mathcal{OCBA}^{\mathbf{EA}}_{\mathbf{LL}}(\beta^\star)$

1. Make an initial number of evaluations n_0 of each individual without evaluations so far. Estimate the ranks by ordering the individuals based on the observed mean values.

2. Determine $\mathcal{C}$: initialize $\mathcal{C} \leftarrow \emptyset$ and add comparisons from the operators as given in Sections 3.3–3.4.

3. WHILE the observed results are not sufficiently sure, i.e. $\mathrm{PBG}_{\mathrm{Slep},\delta^*} > \alpha^\star$ or $\mathrm{EOC}^{\mathrm{Gen}}_{\mathrm{Bonf}} > \beta^\star$, DO

   (a) Allocate new evaluations to the individuals with the OCBA-allocation rule. Individuals leaving the current population $\mathcal{P}$ due to new observations are replaced by the entering individuals. This way comparisons based on position in $\mathcal{P}$ (like tournament selection) others than the one directly affected are not changed.

(b) If ranks have changed from the previous iteration of the ranking procedure, update $\mathcal{C}$: initialize $\mathcal{C} \leftarrow \emptyset$ and add comparisons from the operators as given in Sections 3.3–3.4.

$\mathcal{OCBA}^{\mathrm{EA}}_{\delta^*}$ is called in every iteration of the EA after the offspring has been generated and before replacement. Then, the EA proceeds simply using the ordering given by the sample means.

The number of samples taken by $\mathcal{OCBA}^{\mathrm{EA}}_{\delta^*}$ depends on the problem configuration and the settings of α^* and δ^*. It may be useful to vary α^* over time, as higher accuracy may be needed towards the early and late phases of the algorithm (cf. [Branke 2001]).

3.6 Empirical Evaluation

We empirically compare different sampling mechanisms based on their ability to minimize "incorrect iterations" of the EA. The proposed integration of ranking and selection with EA needs a more exhaustive evaluation in future work.

We stochastically generated $k = 10$ individuals with means distributed according to the negative of an exponential distribution with mean 1 and variances distributed according to an inverted gamma distribution with $\alpha = 100$ and $\beta = 99$ (Figure 3.5). This corresponds to the RPI2 configuration in Chapter 2. Such a distribution with more good than bad individuals seems common in an EA run, at least towards the end of the run, when the algorithm produces many solutions close to the optimum, and a few outliers. To achieve statistically significant results we averaged over 100,000 such randomly sampled populations. This is the same number of macroreplications like in 2.

We compare the frequentist probability of a good generation $\mathrm{PGG}_{\mathrm{iz},\delta^*}$ depending on the expected number of evaluations used by different procedures. To calculate $\mathrm{PGG}_{\mathrm{iz},\delta^*}$, we run the sampling mechanism and look at the resulting order according to sample means. If all decisions required by the scenario (i.e., those defined by $\mathcal{C}$) have been identified correctly taking

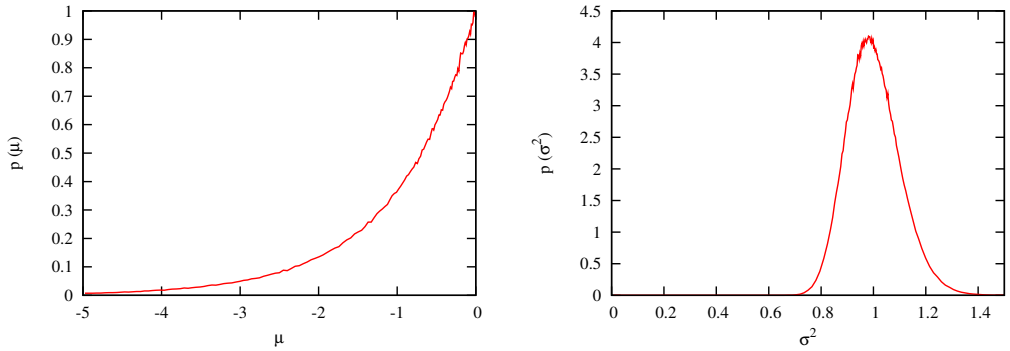

Figure 3.5: Empirical distribution of means (left) and variances (right) for the empirical tests.

into account the indifference zone, the run is counted as successful. Otherwise, it is not successful. $\mathrm{PGG}_{\mathrm{iz},\delta^*}$ is the percentage of correct runs. The parameters α^* and δ^* not only are determinants of $\mathrm{PGG}_{\mathrm{iz},\delta^*}$, but also of the expected total number of samples, $\mathrm{E}[N]$, for a given numerical experiment.

The sampling mechanisms considered are:

1. A standard allocation scheme which samples all individuals equally often. This is denoted by Equal.
2. $\mathcal{OCBA}^{\mathrm{EA}}_{\delta^*}$ for a steady-state EA with population size $\mu = 9$ and one offspring generated
3. $\mathcal{OCBA}^{\mathrm{EA}}_{\delta^*}$ for an evolution strategy with $(5, 10)$ replacement.
4. $\mathcal{OCBA}_{\delta^*}$ to select the best of the 10 individuals.
5. $\mathcal{OCBA}_{\delta^*}$ to give a complete ranking of the 10 individuals.

For all tests, an indifference zone of $\delta^* = 0.2$ is used, i.e. the ordering of a pair of individuals is accepted as correct if the higher ranked individual is not more than 0.2 worse than the lower ranked individual. We use the stopping rule $\mathrm{PGG}_{\mathrm{Slep},\delta^*} > 1 - \alpha^*$, where α^* is varied to generate lines in Figure 3.6.

A complete ranking of the individuals is the most challenging task and requires the highest number of samples to reduce the error $1 - \mathrm{PGG}_{\mathrm{iz},\delta^*}$. The curves for steady-state EA and (5,10) ES are significantly below the

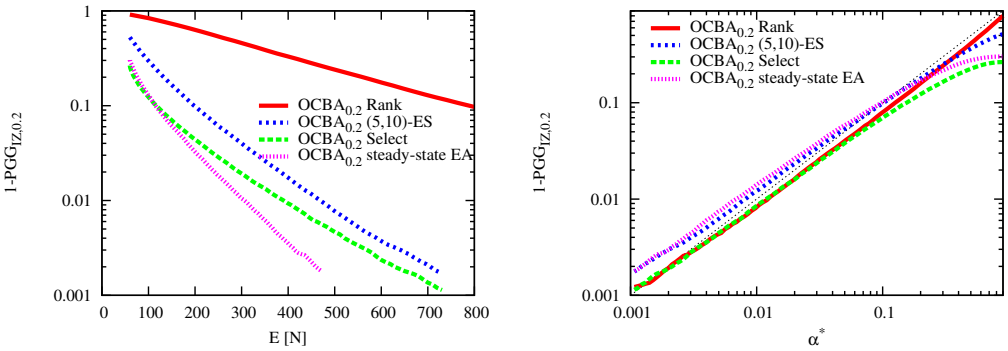

Figure 3.6: Efficiency (left) and target (right) of different sampling mechanisms (expected number of samples required, E[N] and stopping parameter α^*, to reach a certain level of $\mathrm{PGG}_{\mathrm{iZ},\delta^*}$).

curve for the complete ranking, which shows that an EA indeed requires only partial information, and that a lot of samples can be saved by generating only the required information. Interestingly, the steady-state EA operation even requires less samples than identifying only the best individual ($\mathcal{OCBA}_{0.2}$ Select). This is due to the fact that we generate the means according to a negative exponential distribution, i.e. there are several good but few bad individuals, and thus it is relatively easier to identify the worst individual and the better of two random pairs for the steady-state EA than it is to identify the best individual.

Figure 3.7 compares our new $\mathcal{OCBA}$-based EA with standard EAs using the same number of samples for each individual. The $\mathcal{OCBA}$-based sampling allocation schemes are much more efficient than the corresponding Equal allocation variants, which shows that integrating a statistical ranking and selection procedure is beneficial.

For example, to reduce the probability of an erroneous generation, $1 - \mathrm{PGG}_{\mathrm{iZ},\delta^*}$, to 0.02, the standard $(5, 10)$-ES would require an average of 1160 evaluations per generation. Our $\mathcal{OCBA}$-based ES achieves the same accuracy with 385 evaluations per generation. For the steady-state EA, the differences are even larger: the standard steady-state EA would require an average of 845 evaluations per generation. while our $\mathcal{OCBA}$-based EA only requires 240 evaluations. That is, our method saves 67-71% of the samples, and the

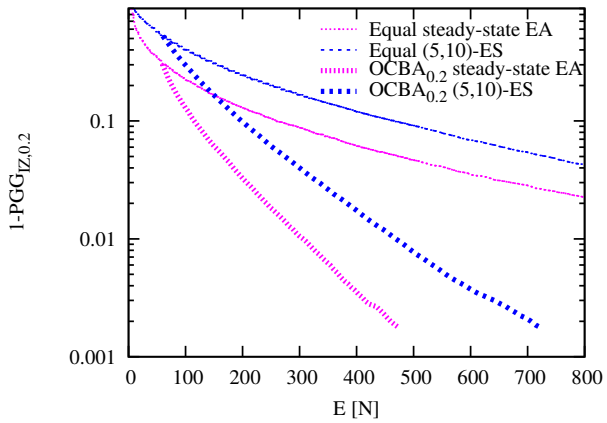

Figure 3.7: Comparison of the new $\mathcal{OCBA}$-based EA (bold lines) and the standard EA with Equal sampling (thin lines).

benefit of our newly proposed methods increases with an increasing desired $\mathrm{PGG}_{\mathrm{iZ},\delta^*}$.

The controllability, i.e. the ability to achieve a given level of $\mathrm{PGG}_{\mathrm{iZ},\delta^*}$ is very high, whereas the target cannot be guaranteed. This holds for all procedures with respect to the chosen configuration.

While this only looks at a single example of an artificially generated iteration, it is to be expected that the benefit of our proposed method will be even larger for larger populations (because the required information becomes an even smaller portion of the information required for a full ranking) and in an iterative EA setting (because $\mathcal{OCBA}$ will be able to re-use the samples of surviving individuals, and those are the individuals that were allocated relatively many samples).

To improve efficiency, prior information about the distribution of means and variances can be acquired from the previous iteration and integrated into the estimation as shown in Section 2.4.

3.7 Setting the Error Probability

With the procedures derived in the previous section an efficient and controllable procedure is available for use in EA. The open question is how to set

the indifference zone δ^* and the error probability $\alpha^\star$ or the loss $\beta^\star$. This setting depends on the use of the EA in the stochastic environment.

1. **Generator**
   The EA is used as generator for solutions, where every solution visited is stored in a database and the best solution is determined after the "optimization" process via selection of the best out of the database. The selection step might be challenging, as a large number of alternatives need to be compared. The advantage of this approach is the guarantee, that the best out of the visited solutions is returned. The disadvantage is that during the search only uncertain information on the current best solution can be given. The ranking procedure only needs to ensure that the optimization does not turn into a random search. This approach was chosen by [Boesel 1999] and is not further investigated in this thesis.

2. **Optimizer**
   The EA is used analogously to the use in a deterministic environment, where the decision maker expects the best solution in the latest population of a run. If the replacement operator ensures that the best (or perceived best) solution survives to the next generation with a high probability, then the returned solution can be determined from the latest population without spending much effort to find the solution from a large database of potential solutions, but instead from a small population.

EA that can be used for the second purpose are generational EA with elitism, steady-state EA and Plus-ES. The latter two replace the worst individual(s), so if the probability for correct generation is $\mathrm{PGG}_{\mathrm{iZ},\delta^*}$, then the probability denoted by $\mathrm{PGS}_{\mathrm{iZ},\delta^*}$ that the best individual survives to the next generation is significantly larger than $\mathrm{PGG}_{\mathrm{iZ},\delta^*}$.

Note that the probability for the best individual to survive t generations is larger than $(\mathrm{PGS}_{\mathrm{iZ},\delta^*})^t$, because the $\mathrm{PGS}_{\mathrm{iZ},\delta^*}$ of successive generations are positively correlated.

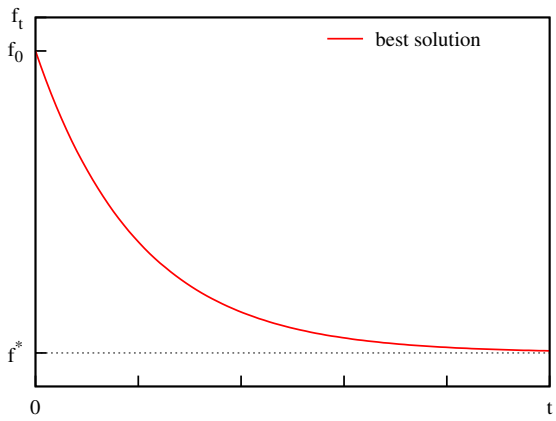

Figure 3.8: Typical convergence curve of an EA depending on number of iterations t.

If the EA is used interactively, it might be necessary to keep the decision maker informed on the currently best solution. Therefore in each generation – or at least for a given interval – the overall best solution found needs to be available (online). This can be easily integrated into $\mathcal{OCBA}_{\delta^*}^{\mathrm{EA}}$, but will modify the behavior of the EA, as samples are used for information that is not needed within the EA. For non-interactive EA the choice of the best solution in the last population might be delayed for after the optimization (offline).

In general the optimal choice of $\alpha^\star$, δ^* and $\beta^\star$ not only depends on the problem, but also on the parameters of the EA like population size, selection pressure, mutation rate and others. Even worse, the optimal setting may vary over the run of an EA.

In the following we derive theoretical conditions for the optimal setting of $\alpha^\star$ or $\beta^\star$ by making assumptions on the trade-off between the benefit of additional accuracy and the price of achieving additional accuracy.

Generally EA converge exponentially fast to the optimum (see Figure 3.8). The speed of the convergence $d = -\frac{\partial}{\partial t}\log(f_t - f^\star)$ depends on the parameter settings. Assume that the convergence speed is locally constant. The convergence speed, when applied to a certain deterministic problem is denoted d_0. In a stochastic environment, d_0 could only be achieved for an error prob-

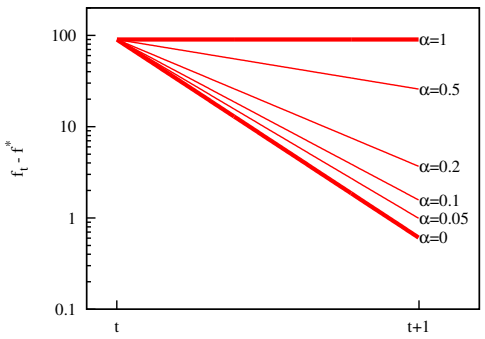

Figure 3.9: Improvement per iteration for different error probabilities α. The slope on this scale corresponds to the convergence speed d_α.

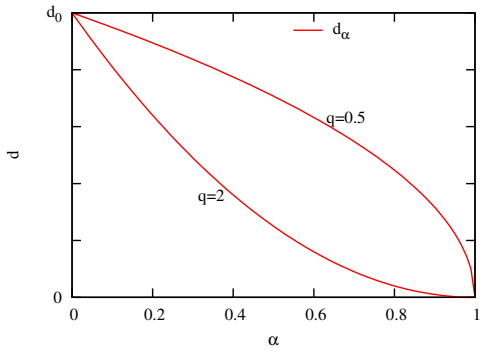

Figure 3.10: Convergence speed d increasing monotonically to the maximal convergence d_0 for decreasing error probability α.

ability $\alpha = 0$, which demands an infinite number of samples n. A higher error probability will result in less selection pressure and therefore the convergence speed is reduced. In the extreme case, if no samples are taken at all, the EA has no information which individuals are better and degrades to a random walk without any improvement ($d_1 = 0$). The convergence speed d_α for the error probability α is therefore in the range $[d_1, d_0]$ (see Figure 3.9). Reasonably, the convergence should increase monotonically with decreasing α. Assume the convergence increases polynomially with exponent q (see Figure 3.10).

$$d_\alpha \quad = \quad d_0(1-\alpha)^q \tag{3.8}$$

With these assumptions, we can give a functional relation between the error probability α in each iteration and the quantitative benefit of additional accuracy for t iterations: $\Delta_t = f_{t_0} - f_{t_0+t} = (f_{t_0} - f^\star)(1 - e^{-d_\alpha t})$. In Figure 3.7 the error probability decreases exponentially in the number of samples per generation n. If no samples are used ($n = 0$), we assume the error probability $\alpha = 1$, so the functional relation is approximately $n = -c \log \alpha$, where c determines the efficiency.

Given a budget of N samples available for optimization, the question arises, how many samples should be spend for the accuracy of a single itera-

tion and how many iterations should be performed. The more iterations the EA runs, the better the result and the more accuracy per iteration, the higher the convergence per iteration. But more accuracy per iteration involves more samples per iteration and therefore budget for less iterations overall.

Taking no samples per iteration allows for an infinite number of iterations, but results in no improvement per iteration and therefore no improvement overall. Using the budget of N samples in one iteration would maximize the improvement in this iteration, but would only allow for one iteration, again resulting in (almost) no improvement. So there is a trade-off between the number of samples per generation and the total number of iterations. Assuming that the number of samples per iteration n is approximately constant, the number of iterations is $T = N/n$. Inserting this together with the relation for samples and error probability into the relation for accuracy gives

$$\begin{aligned} \Delta_T(\alpha) &= (f_{t_0} - f^\star)\left(1 - e^{-d_0(1-\alpha)^q \frac{N}{-c\log\alpha}}\right) \\ &= (f_{t_0} - f^\star)\left(1 - e^{\frac{d_0 N}{c}\frac{(1-\alpha)^q}{\log\alpha}}\right) \end{aligned} \tag{3.9}$$

for a total budget of N samples for the optimization. Maximizing Δ_T is equivalent to minimizing $\frac{(1-\alpha)^q}{\log\alpha}$. It is independent of the deviation between best and current solution $f_{t_0} - f^\star$, the budget N, the maximal convergence speed d_0 and the efficiency of the ranking procedure c. The function has its unique minimum for $q = \frac{1-\alpha}{-\alpha\log\alpha} > 0$. Figure 3.11 shows $\Delta_T(\alpha)$

To choose the error probability α per iteration optimally, the ranking procedure needs to sample until the error probability α holds $\frac{1-\alpha}{-\alpha\log\alpha} < q$. Note that α decreases with additional samples.

However in practice, the convergence exponent q is unknown and depends on the current population. Further research is needed to get a deeper understanding of the relationship between error probability and convergence speed.

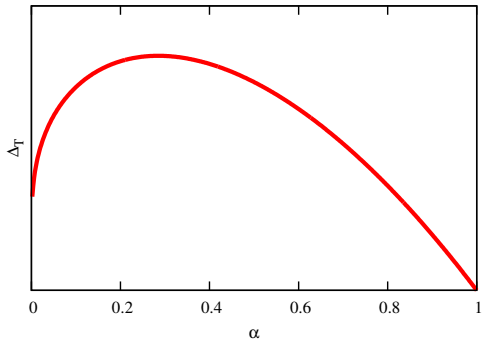

Figure 3.11: Improvement Δ_T for T iterations with varying error probability per iteration α.

3.8 Conclusion

Optimization in noisy environments is challenging, because the noise makes it difficult to decide which of two solutions is better, which is a prerequisite for every optimization algorithm. While noise can usually be reduced by averaging over multiple samples, this is a costly process. In this chapter, we have proposed a new adaptive sample allocation mechanism that attempts to minimize the number of samples required to warrant a proper functioning of an EA. The approach is based on two ideas:

1. Restriction of the focus on those pairwise comparisons that are actually used by the EA. As the empirical results have shown, these comparisons may require less samples than even only identifying the best individual.

2. The use of $\mathcal{OCBA}$, a sample allocation procedure from statistical ranking and selection. This allowed to distribute the additional evaluations to those individuals where they promise the highest benefit, and to stop sampling when there is sufficient evidence for correct selection.

3.9 Future Work

A comprehensive study on the effects of population size, selection pressure and accuracy has to be done for different problems. The current approach

only allows for the imitation of existing methods with higher efficiency. Imitating the behavior of an EA in a deterministic environment results in excessive sampling for typical parameters. Using a fixed sample size scheme as often used in stochastic environments has proven to be suboptimal in Chapter 2.

An approach for choosing a near-optimal error probability is to maximize the improvement per sample Δ_1/n. The improvement per iteration Δ_1 can be estimated by comparing the mean fitness of the old population $\mathcal{P}'$, with the mean fitness of the new parents $\mathcal{M}$. As long as the ratio of estimated improvement and samples does not decrease significantly, the ranking procedure keeps on sampling.

Weighted comparisons, such that the geometric mean over all comparisons holds the target, allows the integration of a higher probability for the best to survive or for more important comparisons. The identification of these will be both promising and challenging.

One further area of future research can be the adaptation of $\mathcal{OCBA}_{\delta^*}^{\mathrm{EA}}$ for multiple objectives, allowing to optimize problems with multiple performance criteria in stochastic environments.

Chapter 4

Using Noise

For noisy optimization problems, there is generally a trade-off between the effort spent to reduce the noise (in order to allow the optimization algorithm to run properly), and the number of solutions evaluated during optimization. However, for stochastic search algorithms like evolutionary optimization, noise is not always bad. On the contrary, in many cases, noise has an effect very similar to the randomness which is purposefully and deliberately introduced e.g. during selection. At the example of stochastic tournament selection, we show that the noise inherent in the optimization problem can be used to partially replace the randomness in selection, thereby reducing the required sampling effort by approximately 50% for typical parameter settings.

4.1 Introduction

Generally, noise is considered harmful, as it may mislead the optimization algorithm. However, noise is not always bad. EA are randomized search algorithms, and most of them use deliberately randomness to purposefully introduce errors into the selection process, primarily in order to get out of local optima.

Therefore, in this chapter we argue that it should be possible to accept the noise inherent in the optimization problem and to use it to (at least partially)

replace the randomness in the optimization algorithm. We propose noise-adjusted tournament selection (NATS) where the probability of accepting the better individual is tuned to reflect the uncertainty in evaluation. As a result, it is possible to dramatically reduce the effect of noise without requiring an excessive number of samples.

Experiments on a simple sphere indicate that the effort in terms of samples for achieving a high selection pressure is not worth the gain in additional solution improvement per generation. It seems that at least for some parameter settings, the best results are obtained with relatively low selection pressure like obtained e.g. with stochastic tournament selection.

4.1.1 Stochastic Tournament Selection

Stochastic tournament selection (STS) has been described in detail in Section 3.4.3. Selecting the better of two individuals with probability $(1-\gamma)$ in a noisy environment can be achieved in two fundamental ways: The standard way would be to eliminate the noise as much as possible by using a large number of samples, and then selecting the better individual with probability $(1-\gamma)$. The noise-adjusted tournament selection proposed here has a different philosophy: instead of eliminating the noise and then artificially introducing randomness, we propose to accept a higher level of noise, and only add a little bit of randomness to obtain the desired behavior.

4.1.2 Basic Notations

Let us denote the two individuals to be compared as i and j. If the fitness is noisy, we assume that the fitness of individual i resp. j is a random variable X_i resp. X_j with $X_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$ resp. $X_j \sim \mathcal{N}(\mu_j, \sigma_j^2)$. However, μ_i and μ_j are unknown, we can only estimate them by sampling each individual's fitness a number of n_i resp. n_j times and using the averages $\bar{x}_i$ and $\bar{x}_j$ as estimators for the fitnesses, and the sample variances s_i^2 and s_j^2 as estimators for the true variances.

If the actual fitness difference between the two individuals is denoted as $\delta = \mu_i - \mu_j$, the observed fitness difference is again a random variable

$D = \bar{x}_i - \bar{x}_j \sim \mathcal{N}(\delta, \sigma_d^2)$. The variance of D depends on the number of samples drawn from each individual, n_i and n_j, and can be calculated as $\sigma_d^2 = \sigma_i^2/n_i + \sigma_j^2/n_j$.

We define $\delta^* = \delta/\sqrt{\sigma_i^2+\sigma_j^2}$ as the standardized fitness difference and $\sigma^* = \sigma_i^2/(\sigma_i^2 + \sigma_j^2)$ as the ratio of the variances. We estimate δ^* by

$$d^* = \frac{\bar{x}_i - \bar{x}_j}{\sqrt{s_i^2 + s_j^2}}$$

The corresponding random variable D^* can be approximated by a non-central t-distribution, where the degrees of freedom ν are given by Welch's approximation. Dividing the nominator and denominator of d^* by $\sqrt{\sigma_i^2 + \sigma_j^2}$ shows that the nominator is a Gaussian random variable with mean δ^* and variance $\frac{\sigma^*}{n_i} + \frac{1-\sigma^*}{n_j}$. The denominator is the root of the sum of two weighted χ^2-distributions. The sum is approximated by the fraction of a χ^2-distribution and its degrees of freedom ν, which are determined by Welch's approximation: $\nu^{-1} = \frac{\sigma^{*2}}{n_i-1} + \frac{(1-\sigma^*)^2}{n_j-1}$. The fraction of a Gaussian random variable with non-zero mean and the root of a χ^2 random variable divided by its degrees of freedom is a Student random variable with the degrees of freedom from the χ^2 random variable.

$$D^* \approx t_\nu\left(\delta^*, \tfrac{1}{\eta}\right) \text{ with } \eta^{-1} = \tfrac{\sigma^*}{n_i} + \tfrac{1-\sigma^*}{n_j} \tag{4.1}$$

The degrees of freedom ν range in $[\min\{n_i - 1, n_j - 1\}, n_i + n_j - 2]$. They are minimal for the extreme values 0 and 1 of σ^* and reach a maximum for $\sigma^* = \frac{n_i-1}{n_i+n_j-2}$. The effective sample size η lies in $[n_i, n_j]$. For equal allocation of n samples among the individuals i and j, the effective sample size η is $n/2$, independent of the variance ratio σ^* and for similar variances ($\sigma^* \approx 0.5$) the degrees of freedom are $\nu \approx n - 2$.

The important conclusion from Equation 4.1 is that the distribution of d^* primarily depends on δ^* (besides n_i and n_j), a property which we will be using later. Also, the above considerations show that d^* is an asymptotically unbiased estimator for δ^*.

For stochastic tournament selection, if $1-\gamma$ is the desired selection probability for the truly better individual, the *desired* selection probability for individual i is

$$\pi_i(\delta^*) = \begin{cases} \delta^* \geq 0 & : \quad 1-\gamma \\ \delta^* < 0 & : \quad \gamma \end{cases} .$$

For Simulated Annealing the *desired* selection probability not only depends on the sign of δ^*, but also on the value. The Metropolis acceptance criterion at the temperature level θ defines the selection probability

$$\pi_i(\delta^*) = \begin{cases} \delta^* \geq 0 & : \quad 1 \\ \delta^* < 0 & : \quad e^{\delta^*/\theta} \end{cases} .$$

We denote with $\hat{\pi}_i(d^*)$ the *implemented* probability for choosing individual i based on the estimated standardized fitness difference d^*, and with $p_i(\delta^*) = \mathbf{E}\left[\hat{\pi}_i(D^*)\right]$ the *actual* selection probability for individual i given a true standardized fitness difference of δ^*. As shown above D^* depends besides δ^* on n_i, n_j and weakly on σ^*.

4.2 The Effect of Noise

4.2.1 Standard Approach

The simplest (and standard) way to apply stochastic tournament selection or simulated annealing would be to ignore the uncertainty in evaluation by making the following assumption:

Assumption 1: The observed fitness difference is equal to the actual fitness difference, i.e. $\delta^* = d^*$.

Individual i is selected with probability $1-\gamma$, if $d^* \geq 0$ and with probability γ, if $d^* < 0$.

$$\hat{\pi}_i(d^*) = \begin{cases} d^* \geq 0 & : \quad 1-\gamma \\ d^* < 0 & : \quad \gamma \end{cases}$$

However, there can be two sources of error: Either we observe a fitness difference $d^* > 0$ when actually $\delta^* < 0$, or vice versa. The corresponding

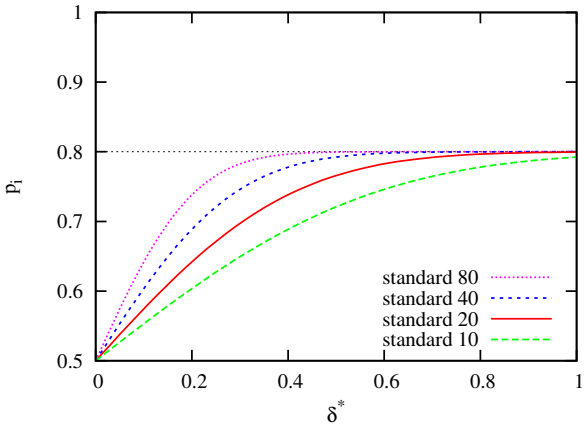

Figure 4.1: True selection probability of individual i depending on the actual standardized fitness difference δ^* for different sample sizes n and equal variances. The dotted horizontal line represents the desired selection probability $(1-\gamma)$.

error probability α can be calculated as

$$\begin{aligned}\alpha &= \begin{cases} \delta \leq 0 \;:\; \mathbf{P}\left(D > 0\right) = 1 - \Phi\left(\frac{-\delta}{\sigma_d}\right) = \Phi\left(\frac{\delta}{\sigma_d}\right) \\ \delta > 0 \;:\; \mathbf{P}\left(D < 0\right) = \Phi\left(\frac{-\delta}{\sigma_d}\right) \end{cases} \\ &= \Phi\left(\frac{-|\delta|}{\sigma_d}\right) = \Phi\left(-|\delta^*\sqrt{\eta}|\right). \end{aligned} \tag{4.2}$$

The overall selection probability for individual i is composed of the probability to observe a difference $d^* > 0$ correctly and then selecting individual i (with probability $1-\gamma$) or else erroneously observe $d^* > 0$ and selecting individual i (with probability γ). Therefore the actual selection probability is $p_i(\delta^*) = (1-\alpha)(1-\gamma) + \alpha\gamma$.

To visualize the effect of the error probability on the actual selection probability p_i, let us consider an example with $\gamma = 0.2$.

Figure 4.1 depicts the resulting true selection probability of individual i depending on the actual standardized fitness difference δ^*. The dotted horizontal line corresponds to the desired behavior in the deterministic case, the bold lines labeled "standard" are the actual selection probabilities due to the noise for sample sizes $n = 80, 40, 20, 10$ for both individuals together, i.e.

$n_i = n_j = 40, 20, 10, 5$ per individual. As can be seen, the actual selection probability for the better individual approaches approximately the desired probability for $\delta^* = 0.35, 0.50, 0.70$ and 1.00, respectively[1]. For $\delta^* \to 0$ it approaches 0.5. The latter fact is unavoidable, since for $\delta^* \to 0$, the signal-to-noise ratio approaches zero, and it becomes basically impossible to determine the better of the two individuals. The interesting question is how quickly p_i approaches $1 - \gamma$, and whether this behavior can be improved. Note that we only show the curves for $\delta^* \geq 0$ (assuming without loss of generality that $\mu_i \geq \mu_j$). For $\delta^* < 0$ the curve would be symmetric to $(0, 0.5)$.

In previous papers, it has been noted that the effect of noise on EA is similar to a smaller selection pressure (e.g. [Miller 1997]). Figure 4.1 demonstrates that there is a notable difference: A lower selection pressure in form of a higher γ would change the level of the dotted line, but it would still be horizontal, i.e. the selection probability for the better individual would be independent of the actual fitness difference. With noise, only the tournaments between individuals of similar fitness are affected. That way, a dependence on the actual fitness values is introduced which somehow contradicts the idea of rank-based selection.

4.2.2 A Simple Correction

If we know that our conclusion about which of the two individuals has a better fitness is prone to some error, it seems straightforward to take this error probability into account when deciding which individual to select. Instead of always selecting the perceived better individual with probability $(1 - \gamma)$, we could try to replace the selection probability of individual i by a function $\hat{\pi}_i(\cdot)$ which depends on the standardized observed fitness difference d^*. Let us make the following assumption:

Assumption 2: It is possible to accurately estimate the error probability α.

Without loss of generality we assume individual i as the better individual.

[1] The deviation at that point falls below 1%. It can be calculated by $-\Phi^{-1}\left(\frac{0.01(1-\gamma)}{1-2\gamma}\right)\sqrt{\eta}$.

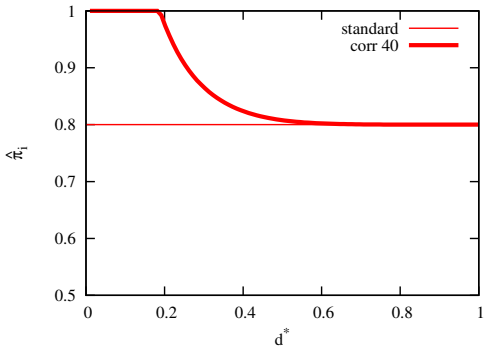

Figure 4.2: Implemented probability $\hat{\pi}_i$ for selecting individual i depending on the observed standardized fitness difference d^* and $n = 40$ samples.

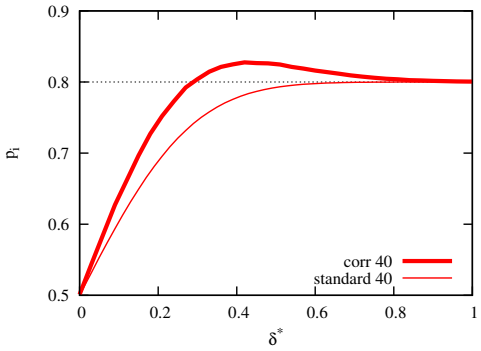

Figure 4.3: True selection probability p_i of individual i depending on the actual standardized fitness difference δ^* and $n = 40$ samples.

It is selected either if we recognize it as the better (probability $(1-\alpha)$) and select the perceived better individual (probability $\hat{\pi}_i$), or if we think it is worse (probability α), but decide to choose the worse individual (probability $1-\hat{\pi}_i$). Thus, since we would like to have an overall true selection probability of $(1-\gamma)$, an appropriate $\hat{\pi}_i$-function could be derived as

$$\begin{aligned}(1-\alpha)\hat{\pi}_i + \alpha(1-\hat{\pi}_i) &\stackrel{!}{=} 1-\gamma \\ \hat{\pi}_i &= \frac{1-\gamma-\alpha}{1-2\alpha}.\end{aligned}$$

$\hat{\pi}_i$ is a probability and can not be smaller than 0, i.e. the above equation assumes $\alpha \leq \gamma < 0.5$. For $\alpha > \gamma$ we set $\hat{\pi}_i = 1$.

Unfortunately, α can not be calculated using Equation 4.2, because we don't know δ^*. It seems straightforward to estimate δ^* by the observed difference d^*. Then, α is estimated as $\hat{\alpha} = \Phi(-|d^*|\sqrt{\eta})$, which is only a biased estimator due to the non-linear transformations. Nevertheless, this may serve as a reasonable first approximation of an optimal $\hat{\pi}_i$-function.

$$\hat{\pi}_i(d^*) = \frac{1-\gamma-\Phi\left(-|d^*|\sqrt{\eta}\right)}{1-2\Phi\left(-|d^*|\sqrt{\eta}\right)} \tag{4.3}$$

Figure 4.2 visualizes this $\hat{\pi}_i$-function (labeled as "corr") for $n = 40$ sam-

ples and $\gamma = 0.2$. As can be seen, the probability to select the better individuals increases when the standardized difference d^* becomes small, and is 1 for $|d^*| < -\Phi^{-1}(\gamma)$ (i.e. the observed better individual is always selected if the observed standardized fitness difference $|d^*|$ is small). Assuming the same parameters as in the example above, the resulting selection probabilities $p_i(\delta^*)$ are depicted in Figure 4.3 (labeled as "corr"). The selection probability approaches the desired selection probability faster than with the standard approach, but then it overshoots before it converges towards $1-\gamma$. Nevertheless, the approximation is already much better than the standard approach (assuming a uniform distribution of δ^*).

4.2.3 Noise Adjusted Tournament Selection

An improvement might by achieved by modifying the function $\hat{\pi}_i(d^*)$ such that the resulting selection probability p_i matches π_i exactly or at least better. An abstract formulation for this is: Given a function $f(.)$ and a random variable Z dependent on x, we look for a function $g(\cdot)$, so that the expected value of the transformed random variable equals the value of the function at the parameter x.

$$\mathbf{E}\left[g(Z_x)\right] = f(x)$$

It is obvious that $g(\cdot)$ equals $f(\cdot)$ if $Z_x = x$, but if the Z_x are randomly disturbed and x is unknown for a given realization z, in general no exact solution exists. $g(\cdot)$ can be chosen to minimize a (weighted) deviation between the expected and the desired value,

$$\min_{g(\cdot)} \int_{-\infty}^{+\infty} w(x)\; d\left(\mathbf{E}\left[g(Z_x)\right], f(x)\right) dx \tag{4.4}$$

for a distance metric $d(\cdot,\cdot)$. The less Z_x varies, the more likely a good fitting function $g(\cdot)$ can be found. Let $p(z,x)$ denote the probability density function (pdf) of Z_x for a given parameter x. If the pdf is n-times differentiable in x, then $\mathbf{E}[g(Z_x)]$ is n-times differentiable in x, too[2], independent of the shape of

[2] $\frac{d^n}{(dx)^n}\mathbf{E}\left[g(Z_x)\right] = \frac{d^n}{(dx)^n}\int_{-\infty}^{\infty} g(z)p(z,x)dz = \int_{-\infty}^{\infty} g(z)\frac{d^n}{(dx)^n}p(z,x)dz$

$g(\cdot)$. This holds for discrete random variables as long as they are continuous in the parameters, too. From this it follows that only continuous $f(\cdot)$ can be matched.

To solve Equation 4.4 we choose the Euclidean distance metric $d(x, y) = \|x - y\|_2$ and $g(\cdot)$ as a stepwise defined function on $m + 1$ intervals $(x_0, x_1]$, ..., (x_m, x_{m+1}), where x_0 and x_{m+1} are $\pm\infty$ and g_j is the value of $g(\cdot)$ in the j-th interval. Further we denote the set of n supportive points by $\tilde{x}_i$. n should be chosen larger than m and some $\tilde{x}_i \notin [x_1, x_m]$ to circumvent boundary effects and force convergence of $g(\cdot)$ towards the boundaries. $p_{ij} = \Pr(Z_{\tilde{x}_i} \in (x_{j-1}, x_j])$ is the probability of observing a value in the j-th interval, if the parameter is $\tilde{x}_i$, w_i is the weight or the loss of a deviation at $\tilde{x}_i$ and $t_i = f(\tilde{x}_i)$ the target value. The mean is then approximated by $\mathbf{E}[g(Z_{\tilde{x}_i})] \approx \sum_{j=1}^{m+1} p_{ij} g_j$. Now Equation 4.4 can be approximated by

$$\min_{g_j, j=0,\ldots,m} \sum_{i=1}^{n} w_i \Big(\sum_{j=1}^{m+1} p_{ij} g_j - t_i\Big)^2 \tag{4.5}$$

This is a quadratic program $\min_{\mathbf{g}} \frac{1}{2}\mathbf{g}^T Q \mathbf{g} + \mathbf{c}^T \mathbf{g}$ with $Q = 2P^T W P$, $\mathbf{c} = -2\mathbf{t}^T W P$, where W is the diagonal matrix of the weights w_i. It can be solved by any quadratic program solver[3].

In the applications below, we interpret $g(\cdot)$ as an applied selection probability, so $g(\cdot) \in [0, 1]$ adds bounds for the variables. For $g(\cdot)$ hitting approximately the target, the variance of the random variable Z_x should be low. If Z_x is the result of a sampling process – as it is below – the variance can be reduced by additional samples.

To apply these results we set $x = \delta^*$, $Z_x = D^*$, $f(\cdot) = \pi_i(\cdot)$ and receive $\hat{\pi}_i(\cdot) = g(\cdot)$. An example curve of $\hat{\pi}_i(\cdot)$ is shown in Figure 4.4.

The calculation of $\hat{\pi}_i$ for 2000 supportive points takes roughly 90 seconds on a 3 GHz PC. The efficiency can be improved by integrating the condition that $g(\cdot)$ has to be point symmetric: $g(x) = 1 - g(-x)$. For calculation of the cdf of the noncentral t-distribution the algorithm found in [Krishnamoorthy and Benton 2003] was used.

[3]We used the freely available OOQP, described in detail in [Gertz and Wright 2003].

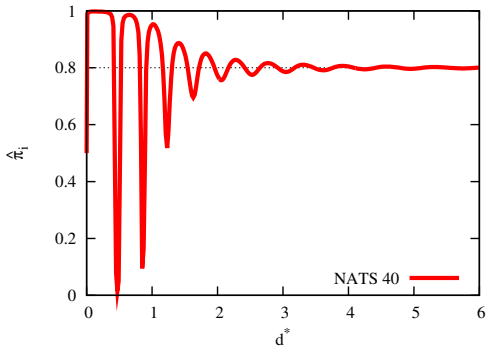

Figure 4.4: Applied selection probability of individual i depending on the observed standardized fitness difference d^* for $\gamma = 0.2$, $n = 20$ and equal allocation. 2000 supportive points are equally spaced between $[-15, 15]$ and $\hat{\pi}_i(\cdot)$ consists of 1000 intervals in $[-10, 10]$.

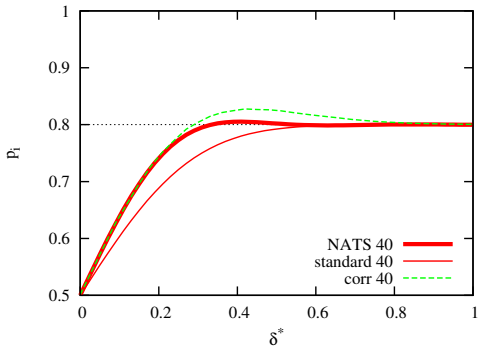

Figure 4.5: True selection probability p_i of individual i depending on the actual standardized fitness difference δ^* and $n = 20$ samples.

Tournament selection based on this optimized acceptance function is denoted noise-adjusted tournament selection (NATS). Note that once the appropriate acceptance function $\hat{\pi}_i(\cdot)$ has been determined, NATS is computationally just as efficient as standard STS. And since the acceptance function primarily depends on γ and the sample size n, typical acceptance functions could be provided by a public repository, making NATS almost as simple to use as STS.

The resulting acceptance function is depicted in Figure 4.4. At first sight, the strong fluctuations seem surprising. However, a steeper ascent of the selection probability can only be achieved by keeping $\hat{\pi}_i = 1$ for as long as possible. The resulting overshoot then has to be compensated by a very low $\hat{\pi}_i$ etc. such that in the end, an oscillating acceptance pattern emerges as optimal. For observed differences up to approximately 0.35, the better individual is always selected, while for a difference around 0.47 the worse individual is selected with high probability. The probabilities vary largely until they converge to the desired probability.

As can be seen in Figure 4.5, despite the oscillating acceptance function

$\hat{\pi}_i(\cdot)$, the curve of achieved selection probabilities is very smooth, and much closer to the actually desired selection probability of γ resp. $(1-\gamma)$ than either the standard approach of ignoring the noise, or the first approximation of an appropriate acceptance function presented in the previous section. Similar results can be obtained for Simulated Annealing, which is not further discussed here.

4.2.4 Sequential Sampling

The sequential selection procedures presented in Chapter 2 showed a large saving compared to the simple equal sampling scheme with a fixed budget of samples. The procedures are designed for the selection of one out of many systems. Regarding a single tournament, only the better of two individuals has to be determined. Although in general, individuals participate in several tournaments, we prescind from that and treat the tournaments as being independent.

Representative for other sequential selection procedures we use Kim and Nelson's procedure to identify the better of two systems and compare it to a new sequential procedure specialized on the comparison of just two systems. The procedure denoted with α_c is based on the assumption that the error probability on each stage should be a constant α_c. Note that because the procedure may stop on any stage due to a large observed fitness difference, the underlying probability distributions are truncated and become rather skewed with an increasing number of iterations. Therefore, we determine the thresholds ϵ_k and the total actual error probability (which also depends on the probability to make a decision in stage k) empirically by simulation. This method is described as Algorithm 1, where $\Phi_{\delta^*,k}(x)$ denotes the cumulative probability distribution for the observed standardized fitness difference in stage k, based on a true fitness difference δ^*. The total number of stages is limited to K.

We use binary search to find a per-stage error probability α_c which yields a total error probability of α at the reference value δ_r^*. Note that the procedure to determine the ϵ_k is computationally more demanding than Kim and

Algorithm 1 Determining thresholds and total error probability

Input: constant error per stage α_c

Generate a large number of pairs of individuals with standardized fitness difference δ_r^*

Sample each individual $n_0 = 4$ times

Estimate $\Phi_{\delta_r^*,0}$ from the observed d^* of all pairs of individuals

FOR $k = 0$ TO $K - 1$ DO {

$\epsilon_k = \max\{0, -\Phi^{-1}_{\delta_r^*,k}(\alpha_c)\}$

Determine probability $p_k^{\delta_r^*}$ to go to next stage based on ϵ_k and $\Phi_{\delta_r^*,k}$

Estimate $\Phi_{\delta_r^*,k+1}$ by truncating at ϵ_k, $-\epsilon_k$, and resampling remaining individuals

}

RETURN total error probability $\alpha = \alpha_c + \sum_{k=1}^{K} \alpha_c \prod_{i=0}^{k-1} p_i^{\delta_r^*}$ and thresholds ϵ_k

Nelson's method. However, this effort has to be spent only once for a given indifference threshold δ_r^* and error probability α, and could in principle be provided by tables.

4.3 Results

4.3.1 Sample size comparisons

Clearly, the methods proposed in the previous section constitute a significant improvement compared to the standard STS. In this section, we try to quantify the improvement in terms of computational effort saved, which is probably the most relevant measure. To this end, we check how many samples are necessary in combination with standard STS to obtain actual selection probabilities similar to when we use the acceptance function generated by our NATS method.

Figure 4.6 shows the actual selection probabilities for standard STS for different number of samples taken (thin lines labeled with 10, 20, 40, 80, 160). On top of that, we have drawn some dashed lines resulting from NATS. As can be seen, the curve obtained from our method combined with 10 sam-

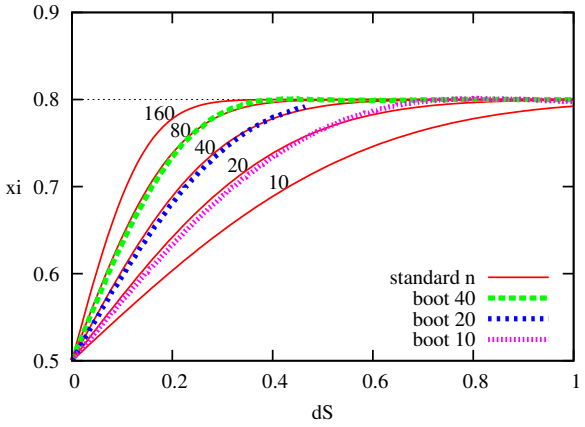

Figure 4.6: Comparison of standard STS and our proposed NATS method for different sample sizes.

ples closely matches the curve obtained with standard STS and 20 samples. These savings seem to be independent of the number of samples: for all curves, NATS allows to cut the number of samples by approximately 50%, still yielding the same accuracy in selection. Since sampling is usually the time-determining factor for complicated real-world optimization problems, this means the run-time of the algorithm is reduced by 50% as well.

4.3.2 Dependency on selection pressure parameter

So far, we have assumed a selection probability for the tournament's better individual of 80%, i.e. $\gamma = 20\%$. As we have explained in Section 3.4.3, the parameter γ allows the algorithm designer to choose the selection pressure. Now we examine how the setting of γ influences the effectiveness our approach.

As we have shown in the previous subsection, with $\gamma = 0.2$, our method allows to save 50% of the evaluations. For $\gamma = 0$, the method is bound to fail, because the underlying trick is to compensate the noise by increasing the probability to select the better individual. Obviously, if the better individual is always selected that idea can not be applied. On the other hand, for $\gamma = 0.5$, we obviously don't need to evaluate at all, but can just select randomly instead. In that case, we could thus save 100% of all evaluations.

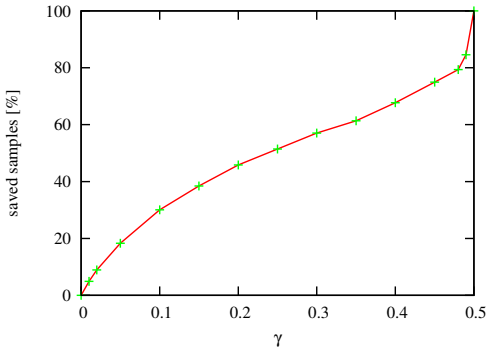

Figure 4.7: Percentage of samples that can be saved by using our proposed noise-adjusted STS, compared to standard STS, without sacrificing selection accuracy.

Figure 4.7 shows the percentage of samples that can be saved using our noise-adapted STS when compared to the standard STS. The savings have been determined similarly to the last subsection, by comparing curves with different sample sizes, and choosing the number of samples n^* for the standard STS such that the resulting selection probabilities p_i are as similar as possible (difference integral) to the noise-adjusted STS with a given number of samples n. The savings are then computed as $1 - n/n^*$. Results are shown for a sample size of $n = 20$, but seem to be independent of the sample size, as the results for sample sizes of 10, 40, or 80 look almost identical (not shown).

As can be seen, in the interesting range of $\gamma \in [0.0, 0.3]$ the savings resulting from the noise-adjusted STS rise quickly with γ, and significant savings can be obtained even for relatively small γ values.

4.3.3 Application to Generational-EA

We present two kinds of experiments. In all cases, we assume the use of stochastic tournament selection, where two individuals are chosen randomly, and the better individual is to be chosen with probability $1 - \gamma = 80\%$. In the first set of experiments, we compare the error probabilities of different methods for a single tournament, depending on the (standardized) true fitness difference. In the second set, we test the approaches on a simple 1000 bit onemax problem, and compare the obtained fitness based on

n	1	10	20	30	40	50
$P(S)$	0.582	0.720	0.765	0.784	0.792	0.796

Table 4.1: Actual selection probability for better individual P(S), depending on the number of samples per individual n and assuming a true standardized fitness difference of 0.35.

the total number of evaluations. For the optimization runs, we assume a simple GA with generational reproduction, population size of 20, one-point crossover with probability 0.6, and bit-flip mutation with mutation probability 1/(chromosome length). Unless stated otherwise, we assume a Gaussian noise with mean 0 and standard deviation $\sigma = 2$.

We compare the different methods based on the *average* population fitness, as the true best is generally unknown in a stochastic environment. Results are averaged over 40 runs.

Selection Error The smallest observable fitness difference possible for the onemax problem is 1 (solutions differ in only 1 bit). Given a Gaussian noise with standard deviation σ, the standardized fitness difference between two such individuals is $\delta^*_{\text{min}} = 1/\sqrt{2\sigma^2}$. If we want to eliminate the effect of noise, we thus require that the selection error is close to zero for $\delta^* > \delta^*_{\text{min}}$.

Let us first consider the case of standard stochastic tournament selection, where the observed better individual is accepted with $(1 - \gamma) = 80\%$ probability. For the assumed Gaussian noise with $\sigma = 2$, we get $\delta^*_{\text{min}} \approx 0.35$. Table 4.1 then shows the actual probability of selecting the truly better individual if the true standardized fitness difference is equal to 0.35, depending on the number of samples used per individual. Any deviation from the desired 80% is due to the noise.

As can be seen, about 40 samples per individual are required to reduce the effect of noise to less than 1% deviation. Figure 4.8 confirms that indeed this is sufficient to eliminate the effect of noise. It shows the average fitness of the population over time, for different levels of noise, assuming 80 samples per tournament (40 per individual). As expected, a noise of $\sigma = 2$ has basically no effect (the lines for $\sigma = 2$ and $\sigma = 0$ are indistinguishable), while larger

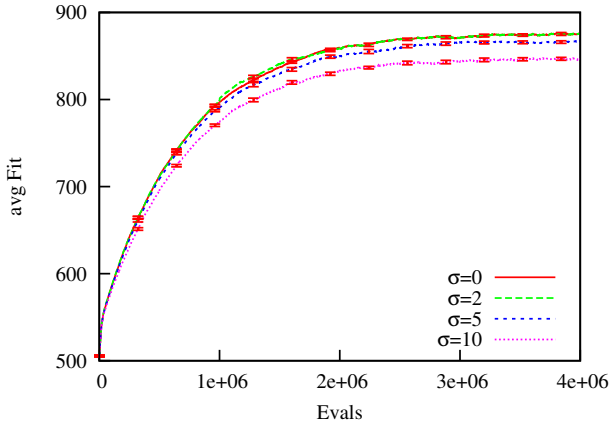

Figure 4.8: Average fitness of population for different levels of noise, with error bars.

noise leads to inferior performance (see $\sigma = 5$ or $\sigma = 10$).

As this example shows, the effect of noise can be eliminated by multiple sampling. However, it also demonstrates the excessive computational burden (80 samples per tournament instead of 2 if the environment would be deterministic). In the following, we show how this computational burden can be reduced dramatically.

Using NATS, a total of 40 samples per tournament (20 samples per individual) should be sufficient to eliminate the effect of noise. For this method, the only relevant issue is when the selection error (probability of observing $d < 0$ although $\delta > 0$ or vice versa) falls below the threshold of $\gamma = 20\%$. Using 40 samples per tournament, this is the case for $\delta_r^* = 0.188$ (see Figure 4.9, line "Standard 40"). Thus, in the following, we are looking for sequential sampling procedures which guarantee an error probability of less than $\alpha = 20\%$ for $\delta^* > \delta_r^* = 0.188$, because these sampling plans should allow us to obtain a selection behavior very close to standard tournament selection based on 80 samples per tournament (and thus complete elimination of noise).

Using Kim and Nelson's procedure and our new proposed sequential sampling procedure with constant α_c, for the given $\delta_r^* = 0.35$ and $\alpha = 0.2$, we observe the error probabilities depicted in Figure 4.9. As can be seen, the

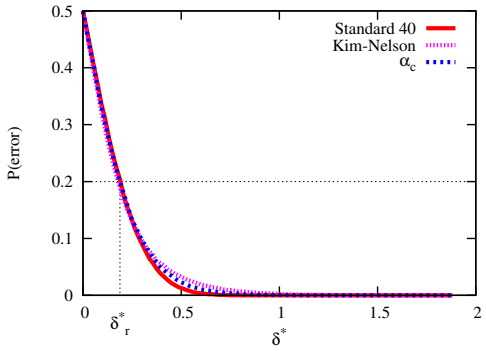

Figure 4.9: Error probability curves for different sampling techniques.

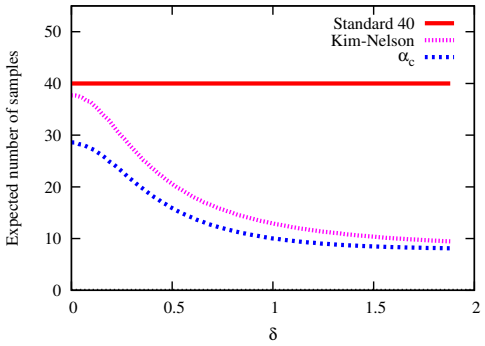

Figure 4.10: Average number of samples taken per tournament for different sampling techniques.

error bounds are observed in all three cases. For $\delta^* > \delta^*_r$, Kim and Nelson's procedure results in a slightly higher error probability than the standard method. Our sampling procedure is better than Kim and Nelson's, but the error is still higher than when the constant sample size is used. This small increase in error might slightly impact the optimization if the sampling method were used in isolation. In our case, however, it is of no importance, because we use it in combination with NATS, which is capable of completely compensating an error of less than $\gamma = 20\%$.

On the other hand, the savings in terms of samples acquired are depicted in Figure 4.10. As can be seen, savings are dramatic for both sequential sampling methods, in particular if the true standardized fitness difference δ^* is large. Our method significantly outperforms Kim and Nelson's procedure, yielding savings of about 14-24% of the samples and lower error probabilities. Therefore in the rest of the paper we only use our sequential sampling method α_c.

As has been stated above, all methods can now be combined with NATS. Figure 4.11 compares the actual selection probability of the better individual depending on the actual standardized fitness difference δ^*, for the following three configurations:

1. The standard stochastic tournament selection with 80 samples per tournament

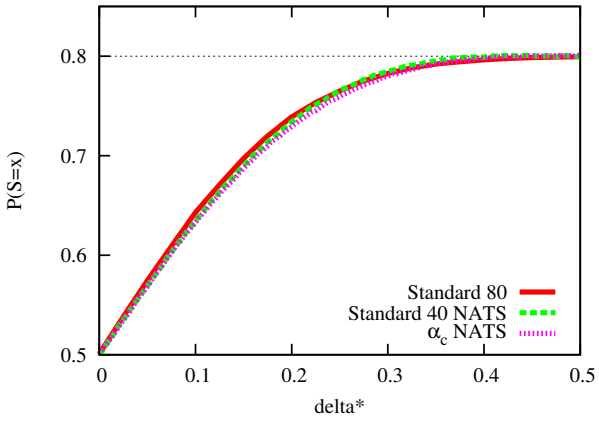

Figure 4.11: Resulting selection probability depending on the true standardized fitness difference δ^* for the standard approach with 80 samples per tournament, the standard approach with 40 samples and NATS, and NATS combined with the α_c sequential sampling scheme.

2. The tournament selection probabilities modified according to NATS and 40 samples per tournament

3. Our sequential sampling procedure with constant error α_c per stage and selection probabilities modified by NATS

As can be seen, all curves look more or less identical. Overall, this means that in order to eliminate the effect of noise, instead of 80 samples per tournament, we only need between 8 and 29 samples on average (depending on the actual fitness difference of the individuals) by using our new sequential sampling mechanism together with NATS.

Behavior during optimization In the previous section, it has been shown that by using appropriate sampling techniques, it is possible to avoid the effect of noise with a much smaller number of samples than if each individual would just be sampled the same number of times. To show that these considerations also hold during optimization, Figure 4.12 empirically compares the convergence curves of the different sampling techniques on the 1000 bit onemax problem. They are all virtually identical to the deterministic case without noise. The corresponding numbers of evaluations per generation are

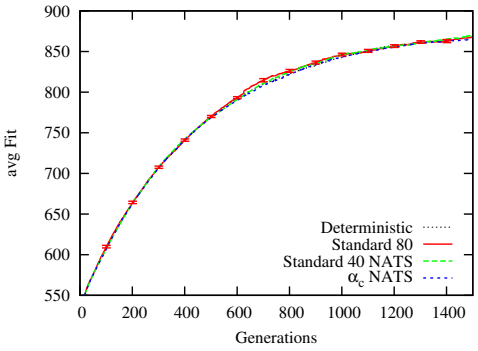

Figure 4.12: Fitness over generations, for different sampling techniques, with error bars. For comparison, the EA's behavior in a deterministic environment is also shown.

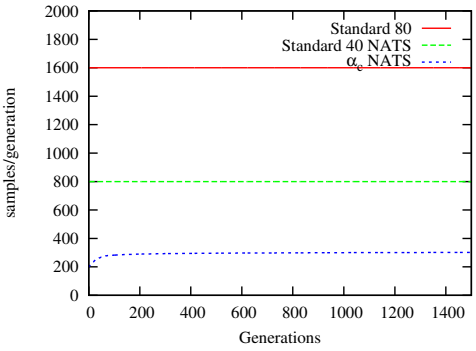

Figure 4.13: Average number of samples required per generation, relative to the deterministic case, for different sampling techniques used.

shown in Figure 4.13 Naturally, the standard approaches sample the same number of times independent of the generation. Surprisingly, the number of samples is also quite constant for the sequential sampling procedure. Only in the very beginning, the number of samples is significantly smaller than later in the run. For our constant α_c method, the sampling effort is never more than 19% of the standard approach, i.e. by integrating the noise into the selection process, we save approximately 81 % of the samples.

4.4 Conclusion and Future Work

In this chapter, we have argued that the noise present in many real-world optimization problems may be used to at least partially replace the randomness in stochastic selection schemes. Instead of trying to estimate an individual's fitness as accurately as possible, a certain level of selection error may be acceptable and can be accounted for simply by adjusting the selection procedure.

At the example of stochastic tournament selection, we have derived two models which determine the selection probability for the better individual depending on the observed fitness difference. The simple model is based on

some simplifying assumptions regarding the distribution of the error probability; the second model is based on a modified acceptance function.

As the results on the onemax problem show, using the proposed noise-adapted tournament selection allows to reduce the number of samples significantly (around 50% for a typical setting) without compromising selection quality. We further reduced the number of samples drastically (up to 81%) by combining NATS with sequential selection procedures.

Future work will include additional improvements resulting from sampling only one individual at a time (instead of both individuals participating in a tournament), and from re-using samples for individuals participating in several tournaments, i.e. to integrate the OCBA procedure derived in Chapter 3.

One further direction of our research is the integration of appropriate population sizing and selection pressure, as it may be crucial to the success of an EA in noisy environments.

Chapter 5

Summary and Outlook

In this thesis, we have developed efficient methods for the application of EA on stochastic problems. To achieve this, we analyzed procedures for statistical selection systematically with respect to different measures and improved them significantly. We showed how to adapt one of the procedures for the needs of EA and identified and modified operators to apply them efficiently in stochastic environments.

For statistical selection we developed a unified experimental setup with standardized configurations allowing for comparison with results from other researchers. Opposed to other publications, where only few values are tabulated, we provide results for a whole range of parameter settings in an easily interpretable graphical form. The new random problem instances model the requirements of selection procedures in practice more realistically. Measurement displays for indifference zone or expected opportunity cost measures are given to compare the procedures with respect to different criteria like efficiency and target.

We identified the expected opportunity cost based procedures $\mathcal{OCBA}_{\mathrm{LL}}(\mathcal{S})$ and $\mathcal{LL}(\mathcal{S})$ to be the most efficient, when a fixed budget of samples is given. In the absence of a budget constraint on average the $\mathcal{OCBA}_{\delta^*}$, $\mathcal{OCBA}_{\mathrm{LL}}$ and $\mathcal{LL}$ procedures provide the best results for a wide variety of configurations, with respect to controllability and efficiency. The procedures based on the indifference zone approach ($\mathcal{KN}$++ and Wald) suffer from the strong depen-

dence on the indifference zone parameter and are hard to control. Nonetheless if a given level of confidence has to be *guaranteed*, $\mathcal{KN}$++ is the procedure of choice. The number of initial samples is a critical parameter for the Bayesian procedures, while $\mathcal{KN}$++ is less sensitive. $\mathcal{KN}$++ allows for correlated output between systems.

From our results we suggest combining the Bayesian allocation procedures with an adaptive stopping rule to substantially improve efficiency compared to the original publications with a fixed budget. The most efficient and controllable stopping rule depends on the desired goal ($\mathrm{EOC}_{\mathrm{Bonf}}$ or $\mathrm{PGS}_{\mathrm{Slep},\delta^*}$). The Bayesian procedures also allow for the incorporation of prior information about problem instances when that is available.

Preliminary work on the analysis and improvement of statistical selection has been presented at the Winter Simulation Conference ([Branke, Chick, and Schmidt 2005a]). More comprehensive results of Chapter 2 have been submitted to Management Science ([Branke, Chick, and Schmidt 2005c]).

Compared to selection, optimization in noisy environments is even more challenging, because not only a few, but combinatorially many alternatives need to be compared. For that the $\mathcal{OCBA}$ procedure is adapted to only include those pairwise comparisons that are actually used by the EA. Most popular selection and replacement operators are covered. We recommend the newly developed latin tournament selection as it combines the low sampling error of stochastic universal sampling with arbitrary selection pressure and its pairwise structure allows for efficient implementation in a statistical procedure. For replacement operators we do recommend those variants that ensure survival of the best known solution, like evolution strategies, steady-state EA or generational EA with elitism, to reduce the effort of identifying the best solution to the last population instead of all visited solutions.

Progress of the optimization is ensured by our adapted statistical selection procedure as it is applied in each generation. The additional samples are distributed to the solutions that promise the highest benefit for the correct functioning of an EA. With the adaptive stopping rules developed, the number of samples automatically varies with the complexity faced in each generation.

The proposed method improves existing approaches for EA on stochastic problems in two aspects: First, it allows to imitate – up to an arbitrary level – the behavior of EA with known fitness values, efficiently and self-adaptive. Second, EA using a fixed number of samples per individual can be improved by allocating samples to individuals, which are more important for the course of an EA than others. A theoretical relationship for the tradeoff between samples used in a single generation and number of generations overall is derived, which could be used to obtain optimal settings for the stopping rule. Unfortunately, the necessary values cannot be determined easily in practice.

Some first results for the integration of statistical selection and EA have been published in [Schmidt, Branke, and Chick 2006].

Tuning the tradeoff between more easily to achieve lower selection pressure and on the other hand less progress of the optimization per generation, often results in relatively low selection pressures compared to those typically used in deterministic environments. Stochastic tournament selection is able to implement such low selection pressures. We derive a method that integrates the remaining noise from the individual's fitness into the probabilistic comparison of stochastic tournament selection by modifying the acceptance function. For typical settings the required number of samples can be reduced by 50%. Combining the method with a sequential comparison procedure allows for a further reduction up to 81% without compromising selection quality.

The general idea for the noise-adjusted tournament selection has been published in [Branke and Schmidt 2003] and received a best paper award. In this thesis an improved approach based on quadratic programming was presented. The combination with sequential comparison has been published in [Branke and Schmidt 2004].

With respect to the achievements obtained in this thesis for the field of evolutionary computation, empirical experiments on a large variety of stochastic problems need to be carried out, in order to realize the improvements. A deeper understanding of the effects of population sizing and selection pressure is needed to tune the parameters of the statistical procedures.

So far, fixing the error probabilities to a certain level combined with high selection pressure, quickly leads to excessive sampling.

Bibliography

Aizawa, A. N. and B. W. Wah (1994). Scheduling of genetic algorithms in a noisy environment. *Evolutionary Computation*, 97–122.

Albert, L. A. and D. E. Goldberg (2001). Efficient evaluation genetic algorithms under integrated fitness functions. Technical Report 2001024, Illinois Genetic Algorithms Laboratory, Urbana-Champaign, USA.

Arnold, D. V. (2002). *Noisy Optimization with Evolution Strategies*. Kluwer.

Arnold, D. V. and H.-G. Beyer (2000a). Efficiency and mutation strength adaptation of the $(\mu/\mu_i, \lambda)$-ES in a noisy environment. Volume 1917 of *LNCS*, pp. 39–48. Springer.

Arnold, D. V. and H.-G. Beyer (2000b). Local performance of the $(\mu/\mu_i, \lambda)$-ES in a noisy environment. In W. Martin and W. Spears (Eds.), *Foundations of Genetic Algorithms*, pp. 127–142. Morgan Kaufmann.

Arnold, D. V. and H.-G. Beyer (2003). A comparison of evolution strategies with other direct search methods in the presence of noise. *Computational Optimization and Applications 24*, 135–159.

Baker, J. E. (1987). Reducing bias and inefficiency in the selection algorithm. In J. Grefenstette (Ed.), *International Conference on Genetic Algorithms*, pp. 14–21. Lawrence Erlbaum Associates.

Baum, C. W. and V. V. Veeravalli (1994). A sequential procedure for multihypothesis testing. *IEEE Transactions on Information Theory 40*(6), 1994–2007.

Bechhofer, R. E., T. J. Santner, and D. M. Goldsman (1995). *Design and Analysis for Statistical Selection, Screening, and Multiple Comparisons*. New York: John Wiley & Sons, Inc.

Benders, J. F. (1962). Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik 4*, 238–252.

Bernardo, J. M. and A. F. M. Smith (1994). *Bayesian Theory*. Chichester, UK: Wiley.

Beyer, H.-G. (1993). Toward a theory of evolution strategies: Some asymptotical results from the $(1 \dot{+} \lambda)$-theory. *Evolutionary Computation 1*(2), 165–188.

Beyer, H.-G. (2000). Evolutionary algorithms in noisy environments: Theoretical issues and guidelines for practice. *Computer methods in applied mechanics and engineering 186*, 239–267.

Birge, J. R. and F. Louveaux (1997). *Introduction to Stochastic Programming*. New York: Springer.

Boesel, J. (1999). *Search and Selection for Large-Scale Stochastic Optimization*. Ph. D. thesis, Northwestern University, Evanston, Illinois, USA.

Boesel, J., B. L. Nelson, and S.-H. Kim (2003a). Using ranking and selection to 'clean up' after simulation optimization. *Operations Research 51*, 814–825.

Boesel, J., B. L. Nelson, and S. H. Kim (2003b). Using ranking and selection to "clean up" after simulation optimization. *Operations Research 51*, 814–825.

Bonn, M., F. Toussaint, and H. Schmeck (2005). JOSCHKA: Job-Scheduling in heterogenen Systemen mit Hilfe von Webservices. In E. Maehle (Ed.), *PARS Workshop Lübeck*, pp. in press. Gesellschaft für Informatik.

Branke, J. (2001). *Evolutionary Optimization in Dynamic Environments*. Kluwer.

Branke, J., S. Chick, and C. Schmidt (2005a, December). New developments in ranking and selection: An empirical comparison of the three main approaches. In M. E. K. et al. (Ed.), *Proc. Winter Simulation Conference*, pp. 708–717. IEEE, Inc.

Branke, J., S. E. Chick, and C. Schmidt (2005b). New developments in ranking and selection, with an empirical comparison of the three main approaches. In M. Kuhl, N. Steiger, F. Armstrong, and J. Joines (Eds.), *Proc. 2005 Winter Simulation Conference*, Piscataway, NJ, pp. in press. IEEE, Inc.

Branke, J., S. E. Chick, and C. Schmidt (2005c). Selecting a selection procedure. Technology and Operations Management Area, INSEAD, Working Paper.

Branke, J. and C. Schmidt (2003). Selection in the presence of noise. In E. Cantu-Paz (Ed.), *Genetic and Evolutionary Computation Conference*, Volume 2723 of *LNCS*, pp. 766–777. Springer.

Branke, J. and C. Schmidt (2004). Sequential sampling in noisy environments. In X. Yao et al. (Ed.), *Parallel Problem Solving from Nature*, Volume 3242 of *LNCS*, pp. 202–211. Springer.

Buchholz, P. and A. Thümmler (2005). Enhancing evolutionary algorithms with statistical sselection procedures for simulation optimization. In N. Kuhl, M. N. Steiger, F. B. Armstrong, and J. A. Joines (Eds.), *Winter Simulation Conference*, pp. 842–852. IEEE.

Butler, J., D. J. Morrice, and P. W. Mullarkey (2001). A multiple attribute utility theory approach to ranking and selection. *Management Science 47*(6), 800–816.

Byers, J. (1991). Basic algorithms for random sampling and treatment randomization. *Computers in Biology and Medicine 112*(21), 69–77.

Cantu-Paz, E. (2004). Adaptive sampling for noisy problems. In K. Deb et al. (Eds.), *Genetic and Evolutionary Computation Conference*, Volume 3102 of *LNCS*, pp. 947–958. Springer.

Chen, C.-H. (1996). A lower bound for the correct subset-selection probability and its application to discrete event simulations. *IEEE Transactions on Automatic Control 41*(8), 1227–1231.

Chen, C.-H., E. Yücesan, L. Dai, and H. Chen (2005). Efficient computation of optimal budget allocation for discrete event simulation experiment. *IIE Transactions*, to appear.

Chen, E. J. and W. D. Kelton (2005). Sequential selection procedures: Using sample means to improve efficiency. *European Journal of Operational Research 166*, 133–153.

Chick, S. E. and K. Inoue (1998). Sequential allocation procedures that reduce risk for multiple comparisons. In D. J. Medeiros, E. J. Watson, M. Manivannan, and J. Carson (Eds.), *Proc. 1998 Winter Simulation Conference*, Piscataway, NJ, pp. 669–676. IEEE, Inc.

Chick, S. E. and K. Inoue (2001a). New procedures for identifying the best simulated system using common random numbers. *Management Science 47*(8), 1133–1149.

Chick, S. E. and K. Inoue (2001b). New two-stage and sequential procedures for selecting the best simulated system. *Operations Research 49*(5), 732–743.

Chick, S. E. and K. Inoue (2002). Corrigendum: New selection procedures. *Operations Research 50*(3), 566.

Chick, S. E. and Y. Wu (2005). Selection procedures with frequentist expected opportunity cost bounds. *Operations Research*, to appear.

Dai, L. (1996). Convergence properties of ordinal comparison in the simulation of discrete event systems. *J. Optimization Theory and Applications 91*(2), 363–388.

de Groot, M. H. (1970). *Optimal Statistical Decisions*. New York: McGraw-Hill.

Di Pietro, A., L. While, and L. Barone (2004). Applying evolutionary algorithms to problems with noisy, time-consuming fitness functions. In *Congress on Evolutionary Computation*, pp. 1254–1261. IEEE.

Evans, M., N. Hastings, and B. Peacock (1993). *Statistical Distributions* (2 ed.). New York: Wiley.

Fitzpatrick, J. M. and J. J. Grefenstette (1988). Genetic algorithms in noisy environments. *Machine Learning 3*, 101–120.

Fu, M. C. (2002). Optimization for simulation. *INFORMS Journal on Computing 14*(3), 192–215.

Fu, M. C., J.-Q. Hu, C.-H. Chen, and X. Xiong (2005). Simulation allocation for determining the best design in the presence of correlated sampling. *INFORMS Journal on Computing*, to appear.

Genz, A. and F. Bretz (2002). Comparison of methods for the computation of multivariate t probabilities. *Journal of Computational and Graphical Statistics 11*(4), 950–971.

Gertz, E. M. and S. J. Wright (2003). Object-oriented software for quadratic programming. *ACM Transactions on Mathematical Software 29*, 58–81.

Goldberg, D., G. Korb, and K. Deb (1989). Messy genetic algorithms: Motivation, analysis , and first results. *Complex Systems 3*, 493–530.

Goldberg, D. E. and K. Deb (1991). A comparative analysis of selection schemes used in genetic algorithms. In G. Rawlins (Ed.), *Foundations of Genetic Algorithms*, San Mateo, CA, USA. Morgan Kaufmann.

Goldsman, D., S.-H. Kim, W. S. Marshall, and B. L. Nelson (2002). Ranking and selection for steady-state simulation: Procedures and perspectives. *INFORMS Journal on Computing 14*(1), 2–19.

Gupta, S. S. and K. J. Miescke (1994). Bayesian look ahead one stage sampling allocations for selecting the largest normal mean. *Statistical Papers 35*, 169–177.

Hammel, U. and T. Bäck (1994). Evolution strategies on noisy functions, how to improve convergence properties. In Y. Davidor, H. P. Schwefel, and R. Männer (Eds.), *Parallel Problem Solving from Nature*, Volume 866 of *LNCS*. Springer.

He, D., S. E. Chick, and C.-H. Chen (2005). The opportunity cost and OCBA selection procedures in ordinal optimization. *submitted*.

Hedlund, H. E. and M. Mollaghasemi (2001). A genetic algorithm and an indifference-zone ranking and selection framework for simulation optimization. In *Winter Simulation Conference*, pp. 417–421. IEEE.

Ho, Y. C., R. S. Sreenivas, and P. Vakili (1992). Ordinal optimization of deds. *Journal of Discrete Event Dynamic Systems 2*(2), 61–88.

Inoue, K., S. E. Chick, and C.-H. Chen (1999). An empirical evaluation of several methods to select the best system. *ACM TOMACS 9*(4), 381–407.

Jin, Y. and J. Branke (2005). Evolutionary optimization in uncertain environments – a survey. *IEEE Transactions on Evolutionary Computation 9*(3), 303–317.

Kall, P. and S. W. Wallace (1994). *Stochastic Programming*. Chichester, UK: John Wiley and Sons.

Kelton, W. D., R. P. Sadowski, and D. A. Sadowski (1998). *Simulation with Arena*. Boston: McGraw-Hill.

Kim, S.-H. and B. L. Nelson (2001). A fully sequential procedure for indifference-zone selection in simulation. *ACM TOMACS 11*, 251–273.

Kim, S.-H. and B. L. Nelson (2005). Selecting the best system. In S. G. Henderson and B. L. Nelson (Eds.), *Handbook in Operations Research and Management Science: Simulation*. Elsevier. (to appear).

Kleywegt, A. J., A. Shapiro, and T. H. de Mello (2001). The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization 12*(2), 479–502.

Krishnamoorthy, K. and D. Benton (2003). Computing discrete mixtures of continuous distributions: noncentral chisquare, noncentral t and the distribution of the square of the sample multiple correlation coefficient. *Computational Statistics and Data Analysis 43*, 249–267.

Law, A. M. and W. Kelton (1991). *Simulation Modeling and Analysis*. New York: McGraw-Hill.

Law, A. M. and W. D. Kelton (2000). *Simulation Modeling & Analysis* (3rd ed.). New York: McGraw-Hill.

Lerch, M., G. Tischler, J. W. von Gudenberg, W. Hofschuster, and W. Kraemer (2001). The interval library filib++ 2.0 - design, features and sample programs. Preprint 2001/4, University of Wuppertal.

Matsumoto, M. and T. Nishimura (1998). Mersenne twister: A 623-dimensionally equidistributed uniform pseudorandom number generator. *ACM TOMACS 8*(1), 3–30.

Michalewicz, Z. and D. B. Fogel (1999). *How to solve it: Modern Heuristics*. Springer.

Miller, B. L. (1997). *Noise, Sampling, and Efficient Genetic Algorithms*. Ph. D. thesis, Dept. of Computer Science, University of Illinois at Urbana-Champaign. available as TR 97001.

Miller, B. L. and D. E. Goldberg (1996). Genetic algorithms, selection schemes, and the varying effects of noise. *Evolutionary Computation 4*(2), 113–131.

Nelson, B. L. and S. Banerjee (2001). Selecting a good system: Procedures and inference. *IIE Transactions 33*(3), 149–166.

Nelson, B. L. and D. Goldsman (2001). Comparisons with a standard in simulation experiments. *Management Science 47*(3), 449–463.

Niederreiter, H. (1992). Random number generation and quasi-monte carlo methods. *SIAM*.

Rinne, H. (1997). *Taschenbuch der Statistik*. Frankfurt am Main: Harri Deutsch.

Rinott, Y. (1978). On two-stage selection procedures and related probability inequalities. *Communications in Statistics*, 799–811.

Schmidt, C., J. Branke, and S. Chick (2006). Integrating techniques from statistical ranking into evolutionary algorithms. Springer.

Stagge, P. (1998). Averaging efficiently in the presence of noise. In A. E. Eiben, T. Bäck, M. Schoenauer, and H.-P. Schwefel (Eds.), *Parallel Problem Solving from Nature V*, Volume 1498 of *LNCS*, pp. 188–197. Springer.

Tong, Y. L. (1980). *Probability Inequalities in Multivariate Distributions*. New York: Academic.

Wald, A. (1947). *Sequential Analysis*. New York: John Wiley.

Wald, A. (1950). *Statistical Decision Functions*. New York: John Wiley, Inc.

Welch, B. L. (1938). The significance of the difference between two means when the population variances are unequal. *Biometrika 25*, 350–362.

Whitley, D. (1989). The GENITOR algorithm and selection pressure: Why rank-based allocation of reproductive trials is best. In J. D. Schaffer (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms*, San Mateo, CA. Morgan Kaufman.

Wilson, J. R. and A. A. B. Pritsker (1984). Experimental evaluation of variance reduction techniques for queueing simulation using generalized concomitant variables. *Management Science 30*, 1459–1472.

ISBN: 978-3-86644-128-6

www.uvka.de