

ENHANCING EVOLUTIONARY ALGORITHMS WITH STATISTICAL SELECTION PROCEDURES FOR SIMULATION OPTIMIZATION

Peter Buchholz
Axel Thümmler

University of Dortmund
Department of Computer Science
August-Schmidt-Str. 12
44227 Dortmund, GERMANY

ABSTRACT

In this paper, we present an evolution strategy for the optimization of simulation models. Our approach incorporates statistical selection procedures that efficiently select the best individual, where best is defined by the maximum or minimum expected simulation response. We use statistical procedures for the survivor selection during the evolutionary process and for selecting the best individual from a set of candidate best individuals, a so-called elite population, at the end of the evolutionary process. Furthermore, we propose a heuristic selection procedure that reduces a random-size subset, containing the best individual, to at most a predefined size. By means of a stochastic sphere function and a simulation model of a production line, we show that this procedure performs better in terms of number of model evaluations and solution quality than other state-of-the-art statistical selection procedures.

1 INTRODUCTION

Simulation models are often used to improve the design of a system or to improve a running system. Improvements often can be achieved by tuning some of the system parameters such that the problem becomes a multi dimensional optimization problem. In fact, the combination of discrete event simulation and optimization is one of the most important and challenging problems in the simulation area and a large number of approaches exists. For recent overviews see (Fu 2002, Law and McComas 2002).

From an abstract mathematical point of view the simulator represents a function $\phi(x_1, \dots, x_n)$ for some input parameter vector $\mathbf{x} = (x_1, \dots, x_n)$. The optimization goal is to find $\max_{\mathbf{x} \in W} E[\phi(\mathbf{x})]$ where the response $E[\phi(\mathbf{x})]$, also denoted as $f(\mathbf{x})$, is the expectation of $\phi(\mathbf{x})$ and W is a feasible range for the parameters. Throughout this paper we consider maximization, but using a negative goal function, minimization problems can be formulated in exactly the same way. Since ϕ is only implicitly represented as a simu-

lation model or in other words as a *black box*, only those optimization methods can be applied which do not exploit the structure of function f . Due to the stochastic nature of the simulation model $\phi(\mathbf{x})$ can only be observed with some statistical fluctuation but confidence intervals can be computed to obtain some confidence in the result estimators (Law and Kelton 2000). From this viewpoint, simulation optimization is similar to multi-dimensional optimization of stochastic functions. On the other hand, the variance of simulation results can be reduced by running longer simulations. However, in contrast to other applications, the evaluation of simulation models is costly such that efficiency is a major aspect in the combination of simulation and optimization. The straightforward combination of an optimization algorithm with a simulation tool where the optimization algorithm calls the simulation tool to perform simulation runs of a fixed length and collects afterwards the results will usually not result in an efficient realization. The problem is that the acceptable variability of simulation results depends on the state of the optimization approach such that optimization algorithm and simulation model have to be tightly coupled (Fu 2002). Since both, optimization algorithm and simulation model are programs running on a computer, such a tight coupling is in principle possible, but requires some new concepts.

Different optimization techniques have been applied for the optimization of simulation models including the response surface method, random search, Kriging models, and stochastic optimization (Fu 2002; Kemper, Müller, and Thümmler 2005). In this paper we develop an approach using evolutionary algorithms and specifically evolution strategies for optimization (Schwefel 1995). Evolution based approaches like evolution strategies or genetic algorithms are robust optimization procedures for the optimization of complex functions. Recently these algorithms have been applied to stochastic functions/models (Arnold and Beyer 2003, Hedlund and Mollaghasemi 2002, Nissen and Propach 1998, Sriver and Chrissis 2004). However, most times evolutionary algorithms have been studied on

benchmark functions providing artificial landscapes with some added noise (Nissen and Propach 1998). This scenario is often far away from real simulation responses where the function is smooth over large parts of the parameter space, the noise depends on the parameter values and the evaluation of the function is costly. To the best of our knowledge, the above-mentioned integration of optimization algorithm and simulation model has been realized only very rudimentarily for evolutionary algorithms with a few notable exceptions, e.g., (Hedlund and Mollaghasemi 2002).

Evolution strategies are population based algorithms, from a given population the best individuals are selected as parents and from these individuals an offspring is generated which builds the next generation together or without the parents. The choice of the parents, which is done deterministically by choosing the individuals with the best responses, depends in a stochastic environment on the quality of the result estimators. Consequently, the evolution strategy has to select the best individuals from a population with a high probability and with a low effort. This is, of course, a stochastic ranking and selection problem, which is commonly known in discrete event simulation and a large number of approaches exists in this area. Early two-stage procedures are Rinott's (1978) procedure or the procedure from Koenig and Law (1985), heuristic extensions to Rinott's procedure are given in (Chen and Kelton 2000, Chen 2002, Chen and Kelton 2003) and novel approaches combining screening and selection are proposed by Boesel, Nelson, and Kim (2003). A recent overview paper of Swisher, Jacobson, and Yücesan (2003) compares different methods. In combination with evolution strategies, the methods are used to select the parents in each step of the algorithm and to select the final optimum. The goal is to find a good result with a low effort, which is different from the goal to avoid a wrong decision in every single step as it is usually assumed for ranking and selection.

In this paper, we propose an enhanced evolution strategy for the optimization of stochastic models that uses statistical selection procedures (i) for the survivor selection during the evolutionary process and (ii) for selecting the best individual from a set of candidate best individuals at the end of the evolutionary process. We compare the performance of evolution strategies when applying different selection procedures in these two phases. In particular, we investigate the impact of different levels of model noise on the quality of the optimal solutions found by the evolution strategy and the number of model evaluations required. Furthermore, we present results for various sizes of the set with best individuals, the elite population, which is kept during the optimization process. Additionally, we propose a heuristic procedure that reduces the random size of the subset found by the screen-to-the-best procedure of Boesel, Nelson, and Kim (2003) to at most a fixed predefined size. Unfortunately, this procedure cannot guarantee a certain

probability of correct selection but experimental results show that it performs better in terms of number of model evaluations and solution quality than other state-of-the-art statistical selection procedures. This holds especially when system configurations are in the indifference-zone (i.e., the difference between the best and second-best configuration is below some threshold d^*).

The paper is organized as follows. In the following section we give a brief overview of evolution strategies, as used for our problem. Afterwards we introduce our evolutionary optimization strategy for discrete event simulation by first introducing ranking and selection procedures and then integrating these procedures within the optimization algorithm. Section 4 evaluates the quality of different selection strategies by means of two examples. Finally, concluding remarks are given.

2 EVOLUTION STRATEGIES

Evolution strategies (ES) (Schwefel 1995) are a class of optimization algorithms adopting strategies from natural evolution to perform optimization. The algorithms are population based. A population contains a fixed number of individuals where an individual is described by a parameter vector \mathbf{x} and possibly a vector of strategy variables \mathbf{m} . In ES from a parent population of size μ a set of λ offspring is generated. New offspring are generated by randomly modifying individuals from the parent population. This operation is denoted as mutation and is done by adding a normally distributed random variable with mean zero to each factor. In some variants also recombination is used such that the factors of two parent individuals are combined to form a new individual, e.g., by selecting the first half of the factors from the first parent and the second half from the second parent.

An ES is generation based and generations are numbered consecutively. Let $\mathcal{P}^{(t)}$ be the parent population of the t -th generation. One can distinguish between $(\mu+\lambda)$ -ES (plus-ES) and (μ,λ) -ES (comma-ES). In the former case the new parents are selected from the old parents and the offspring. In the latter case, new parents are selected only from the offspring. It is problem dependent which of both strategies is superior. Thus, in every generation t , a set $\mathcal{Q}^{(t)}$ of λ new candidate solutions is created from $\mathcal{P}^{(t)}$ by means of variation operators consisting of recombination and mutation of candidate solutions. The candidate solutions to form the population $\mathcal{P}^{(t+1)}$ of generation $t+1$ are selected on the basis of their individual objective function values, called fitness in ES jargon. The selection is deterministic, by selecting the best individuals. Depending on the selection type, selection can be either from $\mathcal{P}^{(t)} \cup \mathcal{Q}^{(t)}$ or only from $\mathcal{Q}^{(t)}$.

Let $x_{ij} \in \mathbb{R}$, $i = 1, \dots, \lambda$ and $j = 1, \dots, n$, be the j -th factor of the i -th individual of $\mathcal{Q}^{(t)}$ and let $m_{ij} \in \mathbb{R}$ be the corresponding mutation strength which corresponds to the stan-

standard deviation of the normal distribution used for mutation. The mutation of individuals of $\mathcal{Q}^{(t)}$ is performed by first adjusting the mutation strength according to

$$m_{ij} := m_{ij} \cdot e^{u_i/\sqrt{2n+u_i}/\sqrt{2\sqrt{n}}}, \text{ for } i=1, \dots, \lambda \text{ and } j=1, \dots, n, \quad (1)$$

where the u_i and u_{ij} are drawn from a standard normal distribution $N(0,1)$ and then mutating the factor values of each individual with its adjusted mutation strength

$$x_{ij} := x_{ij} + m_{ij} \cdot u_{ij}, \text{ for } i=1, \dots, \lambda \text{ and } j=1, \dots, n, \quad (2)$$

where the u_{ij} are again drawn from a standard normal distribution $N(0,1)$. Eqs. (1) and (2) describe an adaptive mutation strategy where the mutation strength is also modified during the evaluation process. This is usually superior to fixed mutation strength for all individuals.

If the factors have to observe restrictions (e.g. $x_{ij} > 0$), then either a new offspring is only accepted if it observes the restrictions or it is modified to observe the restrictions. The algorithm is driven towards the optimum by selecting good solutions as parents and it explores the search space by randomly modifying the factors of good individuals. Since ES is an iterative algorithm, a stopping criterion has to be defined which has to be based on the available information and therefore cannot use the distance to the true optimum as a criterion. Usually a fixed number of generations or missing progress during several generations are used as stopping criteria.

ES falls into the class of metaheuristic optimization procedures. The advantage of the approach is its flexibility and robustness according to wrong decisions during the optimization process and therefore also its robustness according to noisy functions. However, the disadvantage is often a huge effort, since a large number of function evaluations are usually necessary to reach the optimum. The former aspect makes the approach very suitable for optimization of stochastic simulation models, whereas the latter aspect shows that an efficient realization of the approach is of major importance for the practical use of the technique in combination with simulation. This implies that simulation runs have to be long enough to distinguish good from worse individuals in a population, but they need not provide an exact ranking of all individuals since a small percentage of wrong decisions is tolerated by ES.

3 OPTIMIZING STOCHASTIC MODELS WITH EVOLUTION STRATEGIES

3.1 Statistical Ranking and Selection Procedures

In this section our goal is to select an individual from a population with k individuals that has the largest expected response value. Let X_{ij} represent the j -th simulation response (i.e., the j -th fitness value) of individual i , for $i = 1, \dots, k$ individuals. For fixed i , we will always assume that the simulation responses of individual i , i.e., $X_{i1}, X_{i2},$

X_{i3}, \dots , are independently and normally distributed random variables with (unknown) means $\mu_i = E[X_{ij}]$ and unknown and possibly unequal variances $\sigma_i^2 = \text{Var}[X_{ij}]$. This assumption is usually realistic since the X_{ij} may be obtained from independent replications of a simulation or appropriately defined batch means from a single simulation.

In the following we assume the μ_i to be ordered such that $\mu_1 \leq \mu_2 \leq \dots \leq \mu_k$, so that individual k is the best individual among all individuals. We call the event ‘‘CS’’ of selecting individual k a *correct selection*. Note that it is not possible to develop a selection procedure which is independent of the true mean values μ_i and that can guarantee a certain probability of correct selection (Law and Kelton 2000). However, in practice one might not care if we wrongly choose individual $k-1$ if μ_k and μ_{k-1} are close together, i.e., if $\mu_k - \mu_{k-1} \leq d^*$, where $d^* > 0$ is the smallest difference the experimenter feels is worth detecting. The parameter d^* is called the *indifference-zone parameter*.

A widely applied selection procedure in the statistical literature is Rinott’s (1978) two-stage procedure. Given the indifference-zone parameter d^* this procedure can guarantee to select the best individual with a predefined probability of at least P^* , with $1/k < P^* < 1$, provided $\mu_k - \mu_{k-1} \geq d^*$. Let n_0 be the number of initial replicated simulations of each individual. Then the first-stage sample means $\bar{X}_i(n_0)$ and sample variances $S_i^2(n_0)$ are computed, i.e.,

$$\bar{X}_i(n_0) = \frac{1}{n_0} \sum_{j=1}^{n_0} X_{ij} \quad \text{and} \quad (3)$$

$$S_i^2(n_0) = \frac{1}{n_0 - 1} \sum_{j=1}^{n_0} (X_{ij} - \bar{X}_i(n_0))^2, \text{ for } i=1, \dots, k. \quad (4)$$

Based on the number of initial replications and the sample variances $S_i^2(n_0)$ obtained from the first stage, the number of additional simulation replications for each individual in the second stage is $N_i - n_0$, with

$$N_i = \max \left\{ n_0, \left\lceil \left(\frac{h}{d^*} \right)^2 \cdot S_i^2(n_0) \right\rceil \right\}, \quad (5)$$

where $h = h(k, P^*, n_0)$ is a constant depending on k , P^* , and n_0 which solves Rinott’s (1978) integral (see also tables in Bechhofer, Santner, and Goldsman 1995). From the first-stage and second-stage samples the overall sample mean $\bar{X}_i = \frac{1}{N_i} \sum_{j=1}^{N_i} X_{ij}$ is computed and the individual with the largest overall sample mean is selected as the best.

Note, that Rinott’s procedure is a conservative procedure that obtains the predefined probability of correct selection P^* by assuming the *least favorable configuration (LFC)*, i.e., an arrangement of the μ_i such that $\mu_k - d^* = \mu_{k-1} = \dots = \mu_1$, which rarely occurs. Furthermore, the procedure has no mechanism to consider the sample mean of the responses after the first stage, and therefore cannot eliminate clearly inferior individuals prior to conducting additional sampling which introduces unnecessary effort.

To overcome the deficits of Rinott's procedure Chen and Kelton (2000) proposed an enhanced two-stage selection procedure (ETSS) that uses the sample means $\bar{X}_i(n_0)$ after the first stage to adjust the second-stage sample sizes. In fact, they replaced the constant h in Eq. (5) with

$$h_i = \frac{h \cdot d^*}{\max\{d^*, \bar{X}_{\max}(n_0) - \bar{X}_i(n_0)\}}, \quad (6)$$

where $\bar{X}_{\max}(n_0) = \max_{i=1, \dots, k} \{\bar{X}_i(n_0)\}$ is the mean of the best individual after the first stage. Note, that the adjusted values h_i are random variables (since $\bar{X}_i(n_0)$ are random variables) this procedure is only heuristic and cannot guarantee anymore the predefined probability of correct selection P^* . In fact, the probability of correct selection decreases with decreasing first stage sample sizes n_0 . More conservative adjustments of h are recently proposed (Chen 2002, Chen and Kelton 2003).

A further concern with Rinott's procedure is that the first-stage sample sizes n_0 must be the same for each individual $i = 1, \dots, k$. In an evolutionary algorithm where after each generation a selection is performed some individuals will survive for the next generation, and thus, it is desirable to reuse the simulation responses already computed in past generations as initial samples for subsequent survivor selections. Recently Boesel, Nelson, and Kim (2003) extended Rinott's procedure to allow unequal first-stage sample sizes. They replaced the first-stage sample sizes n_0 in Eq. (5) with the (individual) first-stage sample sizes n_{0i} for each individual $i = 1, \dots, k$. Furthermore, replacing $h = h(k, P^*, n_0)$ in Eq. (5) by $h(2, (P^*)^{1/(k-1)}, \min_{i=1, \dots, k} \{n_{0i}\})$ they show that $P\{\text{CS} \mid \mu_k - \mu_{k-1} \geq d^*\} \geq P^*$ still holds.

To eliminate clearly inferior individuals prior to the selection procedure a subset pre-selection procedure can be applied (Bechhofer, Santner, and Goldsman 1995). In the

INPUT: Probability of correct selection P^* , indifference-zone parameter d^* , set H with k individuals in contention, number of already computed samples $n_{0i}, i=1, \dots, k$.
OUTPUT: Set H with retained individuals.
(1) FOR $i = 1, \dots, k$ DO
(2) Compute sample mean and sample variance
$\bar{X}_i = \frac{1}{n_{0i}} \sum_{j=1}^{n_{0i}} X_{ij}, \quad S_i^2 = \frac{1}{n_{0i}-1} \sum_{j=1}^{n_{0i}} (X_{ij} - \bar{X}_i)^2$
(3) Let t_i be the $(P^*)^{1/(k-1)}$ -quantile of the t-distribution with $n_{0i} - 1$ degrees of freedom.
(4) OD
(5) FOR $i = 1, \dots, k$ and $j = 1, \dots, k$ DO
(6) Compute weight
$W_{ij} = \sqrt{\frac{t_i^2 S_i^2}{n_{0i}} + \frac{t_j^2 S_j^2}{n_{0j}}}$
(7) IF $\bar{X}_i < \bar{X}_j - \max\{0, W_{ij} - d^*\}$ DO Remove ind. i from H .
(8) OD

Figure 1: Extended Screen-to-the-Best Procedure of Boesel, Nelson, and Kim (2003)

INPUT: Approximate probability of correct selection P_{app} , indifference-zone parameter d^* , set H with k individuals in contention, maximal subset size m .
OUTPUT: Set H with at most m retained individuals.
(1) Compute adjusted probability $P^* = P_{\text{app}}^{1/(k-m)}$.
(2) Let $n_0 = \min_{i=1, \dots, k} \{n_{0i}\}$ where n_{0i} is the number of samples already computed for individual i .
(3) WHILE $ H > m$ DO
(4) FOR ALL Individuals $i \in H$ DO
(5) Compute samples $X_{i,1}, \dots, X_{i,n_0}$, if not yet computed
(6) OD
(7) Call screen-to-the-best procedure presented in Figure 1 with set $H, P^*, d^*/2$, and n_{0i} samples.
(8) $n_0 := n_0 + 1$
(9) OD

Figure 2: Iterative Subset Selection Procedure

subset selection formulation the goal is to choose a subset $H \subseteq \{1, 2, \dots, k\}$ such that $P\{k \in H \mid \mu_k - \mu_{k-1} \geq d^*\} \geq P^*$. Unfortunately, the size of subset H obtained by these procedures is a random variable and thus, not under control of the experimenter. Figure 1 shows a subset selection procedure that allows different initial sample sizes n_{0i} for individuals $i = 1, \dots, k$. This procedure requires no additional sampling. In the best case the set H contains one individual when the procedure terminates, however, it may also be possible that no individuals are screened out. The procedure in Figure 1 was proposed by Boesel, Nelson, and Kim (2003) and is denoted as *extended screen-to-the-best* procedure by the authors. Furthermore, the authors proposed a combined screening and selection (CSS) procedure that makes use of the extended screen-to-the-best procedure and Rinott's procedure with unequal first-stage sample sizes n_{0i} .

As a final selection procedure we consider a heuristic extension to the screen-to-the-best procedure that reduces the random size subset H to a subset of predefined maximum size m . Figure 2 presents a pseudo-code algorithm of this iterative subset selection (ISS) procedure. From given numbers of initial samples $n_{0i} \geq 2$ for individuals $i = 1, \dots, k$, the procedure continues sampling until the size of the subset reaches the predefined size m . Since this procedure repeatedly calls the screen-to-the-best procedure without resampling previously computed samples $P\{k \in H \mid \mu_k - \mu_{k-1} \geq d^*\} \geq P^*$ cannot be assured for the final subset H in this heuristic approach. In fact, the conditional probability of selecting a subset that contains the best individual, given it passed a previous subset selection, depends on the outcome of the previous subset selection. Thus, to be conservative we choose $d^*/2$ as indifference-zone parameter when calling the screen-to-the-best procedure. This adaptation will be justified when considering the experiments in Section 4. Furthermore, since the size of the subset will be reduced at least $k - m$ times we propose to use $P^* = P_{\text{app}}^{1/(k-m)}$, where P_{app} is an approximate probability of correct selection to be predefined by the experimenter.

3.2 Combining Evolution Strategies with Statistical Selection Procedures

In this section we present an evolution strategy (ES) that incorporates statistical procedures for the selection of best individuals. Throughout this section we consider the plus-strategy for the ES but it should be noted that the algorithm could be applied for comma-strategies in a similar way. Two different phases where statistical selection occurs in the optimization process are distinguished. With S1 we denote the selection strategy used for the survivor selection during the evolutionary process, which repeatedly selects μ individuals from the $\mu+\lambda$ parents and offspring, and with S2 we denote the selection strategy applied for the final selection at the end of the evolutionary process, which selects the best individual from a population of τ candidate best individuals, denoted the *elite population*.

Figure 3 presents a pseudo-code algorithm for the enhanced $(\mu+\lambda)$ -ES. First of all the parameters of the ES are initialized and a feasible search-space is predefined. Furthermore, the parent population \mathcal{P} and the elite population \mathcal{B} is initialized (see steps (1) to (4) in Figure 3). After initialization, the evolutionary process continues until a certain termination condition holds. The most common termination condition is to stop when a predefined number of generations have been passed. Other termination criteria can be to stop if progress gets sufficiently small, i.e., the evolution strategy terminates if the individual of the elite population with the largest sample mean does not change for at least (say) 10 consecutive generations.

In each generation of the evolutionary process, λ individuals are selected uniformly random from the parent population. Then variational operators are applied to these individuals to create the offspring (see steps (6) to (8) in Figure 3). Note that the mutation operator can produce an individual which is outside the search-space. In this case we repeat the mutation until an offspring is generated that lies in the search-space. After mutation, individuals are evaluated according to selection strategy S1. Note that the selection strategies only define the number of evaluations required for each individual. The selection itself is then performed by choosing the individuals with the largest mean values for the parent population of the next generation (see steps (9) and (10) in Figure 3). To update the elite population \mathcal{B} we use the screen-to-the-best procedure of Figure 1. Recall, that this procedure requires no additional sampling. Thus, the quality of the elite population inherently depends on the quality of the selection procedure S1. Furthermore, we use no indifference-zone, i.e., $d^* = 0$, for this selection since our purpose is to screen out clearly inferior individuals such that the limited size of \mathcal{B} is only allocated to the candidate best individuals. Finally, when the ES stops the best individual is determined according to selection strategy S2.

```

(1) Predefine survivor selection strategy S1, final selection strategy S2, and size of elite population  $\tau$ 
(2)  $t = 0$ 
(3) Initialize parent population  $\mathcal{P}^{(0)}$  with  $\mu$  individuals created uniformly random in the region of interest
(4) Initialize elite population  $\mathcal{B}$  with  $\min\{\mu, \tau\}$  individuals from  $\mathcal{P}^{(0)}$ 
(5) WHILE Termination condition not satisfied DO
(6)   Create offspring population  $\mathcal{Q}^{(t)}$  with  $\lambda$  individuals each selected uniformly random from parent population  $\mathcal{P}^{(t)}$ 
(7)   Adjust mutation strength of individuals in population  $\mathcal{Q}^{(t)}$  according to Eq. (1)
(8)   Mutate individuals in population  $\mathcal{Q}^{(t)}$  according to Eq. (2)
(9)   Evaluate individuals of  $\mathcal{Q}^{(t)} \cup \mathcal{P}^{(t)}$  according to selection strategy S1
(10)  Select  $\mu$  individuals from  $\mathcal{Q}^{(t)} \cup \mathcal{P}^{(t)}$  with largest mean response value for population  $\mathcal{P}^{(t+1)}$ 
(11)  Compute from  $\mathcal{B} \cup \mathcal{Q}^{(t)} \cup \mathcal{P}^{(t)}$  a subset H of random size according to the screen-to-the-best procedure presented in Figure 1 with  $d^* = 0$ .
(12)  Reinitialize  $\mathcal{B}$  with  $\min\{\tau, |H|\}$  individuals from H with largest mean values.
(13)   $t := t+1$ 
(14) OD
(15) Evaluate individuals of  $\mathcal{B}$  according to selection strategy S2.
(16) RETURN Individual with largest mean response value among all individuals in  $\mathcal{B}$ 
    
```

Figure 3: Pseudo-Code of the Enhanced Evolution Strategy for Plus-Selections

4 QUANTITATIVE RESULTS FOR THE ENHANCED EVOLUTION STRATEGY

4.1 Impact of Statistical Selection Procedures

In this section, we show with several experiments the impact of applying statistical selection procedures on the performance of an evolution strategy. As an experimental model we consider an appropriately scaled n -dimensional sphere function

$$f(\mathbf{x}) = 1 - \frac{1}{4n} \mathbf{x}^T \mathbf{x}, \quad (7)$$

which is designed such that $f(\mathbf{x}) \in [0,1]$ for $\mathbf{x} \in [-1,2]^n$ and the maximum of f is at $f(\mathbf{0}) = 1$ and the minimum of f is at $f(2, \dots, 2) = 0$. To consider the stochastic nature of a simulation model we add normally distributed noise to the undisturbed model function f , where the noise strength is determined by a *variance surface* function $g(\mathbf{x})$, i.e.,

$$f_{\text{sim}}(\mathbf{x}) = f(\mathbf{x}) + g(\mathbf{x}) \cdot N(0,1), \quad (8)$$

where $N(0,1)$ has standard normal distribution. In our experiments we consider the variance surface

$$g(\mathbf{x}) = \sigma \cdot \left(1 + \frac{1}{2n} \sum_{i=1}^n \sin(\gamma \pi x_i) \right), \quad (9)$$

which induces varying variances in f_{sim} over the search-space. In particular, Eq. (9) is designed such that the standard deviation of $f_{\text{sim}}(\mathbf{x})$ varies between 0.5σ and 1.5σ for $\mathbf{x} \in \mathbb{R}^n$. The parameter γ is to adjust the frequency of this

variation, i.e., increasing γ increases the fluctuation rate of the standard deviations between 0.5σ and 1.5σ . Note that $\gamma = 0$ induces constant noise with standard deviation σ in f_{sim} . Furthermore, noise at the maximum of f_{sim} is $g(\mathbf{0}) = \sigma$ independent of the choice of γ . In the experiments we denote σ the *model noise strength* and γ the *model noise fluctuation*, respectively. Note, that for a model noise strength of σ , the mean absolute deviation of $f_{\text{sim}}(\mathbf{x})$ from $f(\mathbf{x})$ is $\sigma_{\text{abs}} = \sigma \cdot \sqrt{2/\pi}$. Since the true optimal response is $f(\mathbf{0}) = 1$ we can interpret the noise induced in $f_{\text{sim}}(\mathbf{0})$ due to $g(\mathbf{0})$ to be of $100 \cdot \sigma_{\text{abs}}$ percent. The two-dimensional response surface of $f(\mathbf{x})$ and variance surface $g(\mathbf{x})$ over the search-space $[-1,2]^2$ are shown in Figures 4 and 5, respectively.

We compare the performance of evolution strategies when applying different selection procedures for survivor selection and final selection, respectively (see strategies S1 and S2 in Figure 3). In particular, we compare the combined screening and selection procedure, abbreviated with CSS, the enhanced two-stage selection procedure, abbreviated with ETSS, and the iterative subset selection procedure, abbreviated with ISS, as discussed in Section 3.1. Additionally, we consider the following two simple heuristic selection strategies:

MEAN: Perform exactly n_0 replicated simulations of each individual and choose the one with the largest mean value.

CONF: Perform exactly n_0 replicated simulations of each individual and continue replicating until the width of the $100 \cdot P^*$ percent t-confidence interval becomes smaller than d^* . Then choose the individual with the largest mean value.




Figure 4: Response Surface $f(\mathbf{x})$ in Two Dimensions




Figure 5: Variance Surface $g(\mathbf{x})$ for $\sigma = 0.2$ and $\gamma = 1.0$

In all experiments we apply a (5+5)-ES on the two-dimensional sphere function (7) and choose the individuals for the first generation uniformly random in search-space $[-1,2]^2$. For all selection procedures an initial first-stage sample size $n_0 = 10$ and a probability of correct selection $P^* = 0.9$ is used. Furthermore, we assume the indifference-zone parameter $d^* = 0.1$ for the survivor selections and $d^*/2$ for the final selections in order to be more accurate at the end. To introduce unequal variances across the search-space we choose $\gamma = 1$ if not mentioned otherwise. For performance evaluation of the evolution strategy two performance indices are considered:

- (i) The distance δ between the response of the best solution found by the evolution strategy and the true optimal response, i.e., $\delta = 1 - f(\mathbf{x}_{\text{opt}})$ where \mathbf{x}_{opt} is the result of the evolution strategy, and
- (ii) the overall number of model evaluations required by the evolution strategy until it returns the result \mathbf{x}_{opt} .

In order to obtain the performance indices with high confidence we made 10,000 independent replications of the complete optimization process for each point in Figures 6 to 11.

In a first experiment the progress of the ES with increasing number of generations is investigated when applying different survivor selection strategies (see Figure 6). In this experiment the model noise strength is $\sigma = 0.2$ and the size of the elite population is $|\mathcal{B}| = 1$, i.e., no final selection is used. From Figure 6 we see that for all selection strategies, except the MEAN(10) selection, the distance value




Figure 6: Performance Indices for Increasing Number of Generations and Different Survivor Selection Strategies

approaches 0.01 with increasing number of generations. Note, that this value is about 10 times smaller than the indifference-zone parameter d^* . The reason for this is that the sphere function is very flat, such that the ES produces even after the first generation a result with a distance below 0.1 from the optimal response (see also innermost contour circle around the optimum in Figure 4, which corresponds to a response of 0.9). Considering the number of evaluations required to obtain the results, the selection strategies differ significantly. In fact, the ISS requires about 2/3 of the iterations required by the CSS.

In the following experiments we consider a fixed number of 50 generation of the ES until termination. Figures 7 and 8 show the performance indices for σ increasing from 0.01 to 0.3 when applying different survivor selection strategies. In Figure 7 no final selection is considered whereas in Figure 8 the ISS procedure is used for final selection with an elite population of size 10. Comparing the different survivor selection strategies we first observe that distance from the optimal solution increases linearly when applying the MEAN selection strategy. In fact with MEAN selection the accuracy of the solution cannot be controlled automatically with respect to increasing noise. With the other selection strategies the solution accuracy can be significantly increased for all noise levels, but with the cost of a (automatically controlled) larger number of model replications. Furthermore, with the same number of evaluations the ISS procedure produces much better results than the MEAN value approach, e.g. considering a model noise strength of $\sigma = 0.23$ in Figure 7, where MEAN(50) and ISS require similar number of evaluations but ISS returns a solution with distance 0.011 and MEAN(50) with 0.015, respectively. Comparing the number of evaluations required by the selection strategies we see that the ISS procedure requires the smallest number of evaluations even when producing the best solutions with respect to the distance measure. Recall, that Figure 8 additionally investigates the impact of applying the ISS procedure for final selection. Compared to Figure 7, for all selection strategies the distance measure can be improved about 40% whereas the number of evaluations increases only about 14%. This is a clear argument for using an elite population during the optimization process.

In the next experiments we further investigate the impact of applying an elite population. Figures 9 and 10 show the performance indices when increasing the size of the elite population from 1 to 30 for different final selection strategies. For survivor selection we used the MEAN(10) and ISS procedures, respectively. In this experiment the model noise strength is $\sigma = 0.2$. For all strategies an increase in elite population size decreases the distance measure and increases the number of evaluations, except for MEAN(10). Observe, that the ETSS procedure does not further improve when the elite population size increases




Figure 7: Performance Indices for Increasing Model Noise Strength, Different Survivor Selection Strategies, and no Final Selection




Figure 8: Performance Indices for Increasing Model Noise Strength, Different Survivor Selection Strategies, and ISS Final Selection

above 8. The reason for this is that the computation of second-stage sample sizes is based on accurate sample means of the first stage, which are not given for $n_0 = 10$ samples.

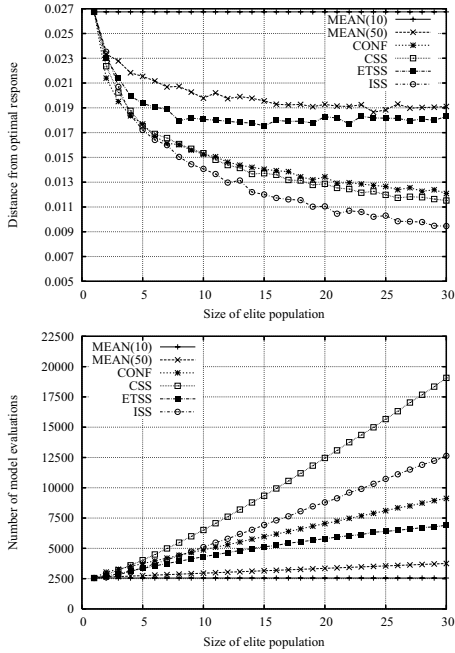


Figure 9: Performance Indices for Increasing Size of Elite Population, Different Final Selection Strategies, and MEAN(10) Survivor Selection

To tackle this problem, Chen (2002) proposed a conservative adjustment to the ETSS procedure. We also checked this procedure and obtained better results, but, as expected, for a larger number of evaluations. In Figures 9 and 10, the ISS procedure produces almost always the best results with respect to the distance measure; however, it also requires a large number of evaluations. Comparing MEAN(10) and ISS survivor selection, we observe an improvement of about 100% in the distance measure with ISS, whereas the number of evaluations increases only by about 10%. Thus, a conclusion from Figures 9 and 10 is that a simple MEAN selection as usually applied in evolutionary algorithms, should not be the method of choice.

In a final experiment in Figure 11, we considered the impact of model noise fluctuation on the performance measures. Again, the model noise strength is $\sigma = 0.2$ and the size of the elite population is $|\mathcal{B}| = 1$. From the curves we see how sensitive the selection procedures are against unequal variances across the search-space. Recall, that $\gamma = 0$ corresponds to constant variances across the search-space. It can be observed that, procedures CONF, CSS, and ETSS are very stable with respect to the distance measure. Nevertheless, the stability requires an increasing number of evaluations. Considering the MEAN(50) and ISS procedure, we observe an increase in the distance measure for γ between 0 and 1. For the ISS procedure the number of evaluations also increases. This is not desirable, but comparing the absolute numbers ISS produces always the best

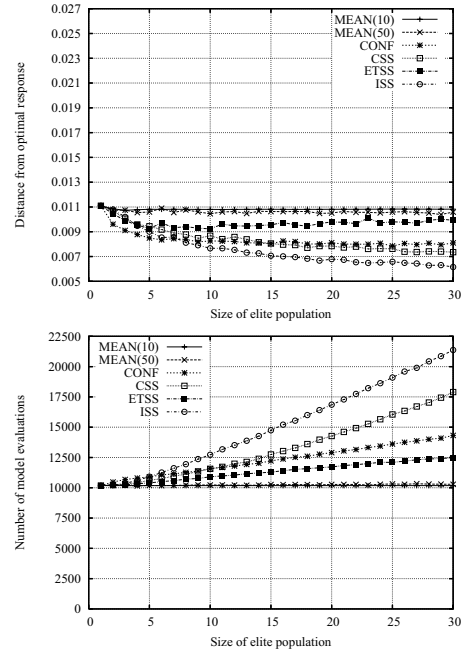


Figure 10: Performance Indices for Increasing Size of Elite Population, Different Final Selection Strategies, and ISS Survivor Selection




Figure 11: Performance Indices for Increasing Noise Fluctuation and Different Survivor Selection Strategies

results with respect to the distance measure as well as the number of evaluations. From the figure we conclude that the performance of ISS as shown in Figures 6 to 10 is further improved when there is small variance fluctuation, i.e., $\gamma < 1$.

Finally, we state two reasons why selection strategies based on Rinott's procedure such as CSS require more evaluations than heuristic procedures like ISS or ETSS: (i) the second-stage sample size is based on the variance estimates of the first n_0 replications, which might give imprecise estimates resulting in a higher second-stage sample size, and (ii) Rinott's procedure does not incorporate mean estimates when computing the second-stage sample size, i.e., it assumes the least favorable configuration.

4.2 Optimization of a Production Line

To provide a more realistic application example, we consider a production line of a manufacturing plant comprising N service queues arranged in a row, that is, parts leaving a queue after service are immediately transferred to the next queue. In this example, we assume all queues to have finite capacity $K = 10$. Arrivals of parts to the first queue occur according to a Poisson process with rate $\lambda = 0.5$. Each queue comprises a single server with first-come, first-served (FCFS) service discipline and exponentially distributed service time. The service rate (i.e., speed of the server) at queue n is denoted with μ_n . Furthermore, the vector $\boldsymbol{\mu} = (\mu_1, \dots, \mu_N)$ of service rates is subject to be optimized by the evolution strategy with respect to a revenue function

$$R(\boldsymbol{\mu}) = \frac{r \cdot X(\boldsymbol{\mu})}{c_0 + \mathbf{c}^T \cdot \boldsymbol{\mu}} - c_1, \quad (10)$$

where $X(\boldsymbol{\mu})$ is the throughput of the production line (i.e., the time-averaged number of parts leaving the last queue) for service rates $\boldsymbol{\mu}$ and r , c_0 , c_1 and \mathbf{c} are constants representing a revenue factor, basic costs that occur independent of service rates and a vector with cost factors for each server, respectively. Since $X(\boldsymbol{\mu})$ is decreasing for decreasing any of the μ_n , $n=1, \dots, N$, the revenue function (10) clearly quantifies the trade-off between a high throughput (i.e., a high production rate) and costs of providing fast service. To compute $X(\boldsymbol{\mu})$ quantitative analysis of the model must be conducted. For general models this can only be done by discrete-event simulation. Nevertheless, for the production line model, also numerical transient analysis of the underlying continuous-time Markov chain can be applied for its quantitative solution (Bause, Buchholz, and Kemper 1998).

In the experiments, we used the production line with $N = 3$ queues and considered transient results at time $t = 1000$ starting with an empty system at time $t = 0$. The revenue factor and the cost factors are assumed to be $r = 10,000$, $c_0 = 1$, $c_1 = 400$ and $\mathbf{c}^T = (1, 5, 9)$. Furthermore, the search-space is bounded to $\boldsymbol{\mu} \in [0, 2]^3$, which results in values of $R(\boldsymbol{\mu})$ ranging approximately between -400 and 98.5 . Since the optimal solution is unknown we cannot use the distance measure (i) as performance index. Instead we consider the numerical solution $R_{\text{num}}(\boldsymbol{\mu}_{\text{opt}})$, where $\boldsymbol{\mu}_{\text{opt}}$ is the

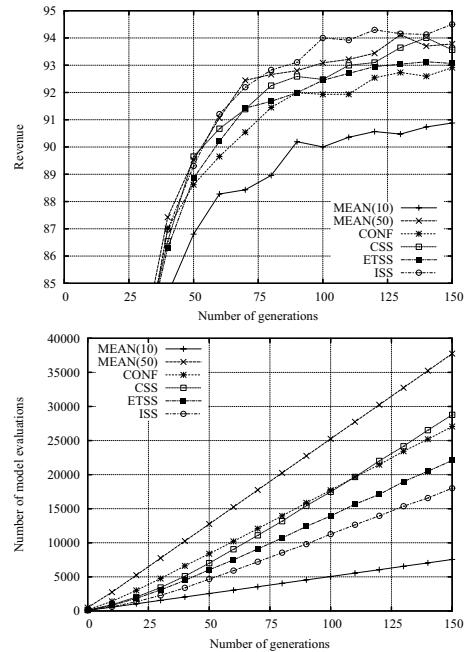


Figure 12: Performance Indices for Increasing Number of Generations and Different Survivor Selection Strategies

result of the evolution strategy. Figure 12 plots the performance indices obtained from 2,500 independent replications of the complete optimization process. In this experiment we considered different survivor selection strategies with $d^* = 10$ and no final selection. It can be observed that ISS produces the best results on average while it requires the smallest number of evaluations. The best solution found during our experiments was $R_{\text{num}}(0.54, 0.45, 0.42) = 98.46$. The average revenue of the solutions found by ISS is 94. Considering the range from -400 to 98.5 of possible revenue values this corresponds to 99% of the best solution.

5 CONCLUSIONS

In this paper, we presented an evolution strategy for the optimization of simulation models. The proposed evolution strategy determines the parent population of consecutive generations using a statistical selection procedure. Furthermore, we propose to use statistical selection at the end of the evolutionary process for selecting the best individual from a set of candidate best individuals collected during the evolutionary process, i.e., the elite population. As a second contribution we propose a heuristic iterative subset selection (ISS) procedure that reduces a random-size subset, containing the best individual, to a maximum predefined size.

By means of two application examples, we investigated the impact of applying statistical procedures for survivor selections and the final selection. The results show, that the performance of the ES is improved when using sta-

tistical selection procedures instead of applying a simple selection based on mean response values. Furthermore, the ISS procedure performs best compared to the other selection strategies for survivor selections as well as the final selection.

In future work, we plan to investigate possible improvement of the evolution strategy when using common random numbers (CRN) across systems for the simulation model and compare the performance with selection procedures especially designed for CRN. A second concern is to study the evolution strategy when using a fixed budget allocation, i.e., the maximal number of model evaluations is a predefined value. Finally, we plan to compare the proposed optimization approach with other optimization approaches, especially those based on response surface methodology (Kemper, Müller, and Thümmeler 2005).

REFERENCES

- Arnold, D.V., and H.G. Beyer, A Comparison of Evolution Strategies with Other Direct Search Methods in the Presence of Noise, *Computational Optimization and Applications* **24**, 135-159, 2003.
- Bause, F., P. Buchholz, and P. Kemper, A Toolbox for Functional and Quantitative Analysis of DEVS, *Proc. 10th Int. Conf. on Modelling Tools and Techniques for Computer and Communication System Performance Evaluation, Palma de Mallorca, Spain, Lecture Notes in Computer Science* **1469**, 356-359, Springer, 1998. <www4.cs.uni-dortmund.de/APNN-TOOLBOX>
- Bechhofer, R.E., T.J. Santner, and D.M. Goldsman, *Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons*, John Wiley & Sons, 1995.
- Boesel, J., B.L. Nelson, and S.H. Kim, Using Ranking and Selection to "Clean up" After Simulation Optimization, *Operations Research* **51**, 814-825, 2003.
- Chen, E.J., A Conservative Adjustment to the ETSS Procedure, *Proc. Winter Simulation Conference, San Diego, CA*, 384-391, 2002.
- Chen, E.J., and W.D. Kelton, An Enhanced Two-Stage Selection Procedure, *Proc. Winter Simulation Conference, Orlando, FL*, 727-735, 2000.
- Chen, E.J., and W.D. Kelton, Inferences From Indifference-Zone Selection Procedures, *Proc. Winter Simulation Conference, New Orleans, LA*, 456-464, 2003.
- Fu, M.C., Optimization for Simulation: Theory vs. Practice, *INFORMS Journal on Computing* **14**, 192-215, 2002.
- Hedlund, H.E., and M. Mollaghasemi, A Genetic Algorithm and an Indifference-zone Ranking and Selection Framework for Simulation Optimization, *Proc. Winter Simulation Conference, Arlington, VA*, 417-421, 2001.
- Kemper, P., D. Müller, and A. Thümmeler, Combining Response Surface Methodology with Numerical Models for Optimization of Class Based Queueing Systems, *Proc. Int. Conf. on Dependable Systems and Networks (DSN), Yokohama, Japan*, 2005.
- Koenig, L.W., and A.M. Law, A Procedure for Selecting a Subset of Size m Containing the l Best of k Independent Normal Populations, With Applications to Simulation, *Communications in Statistics – Simulation and Computation* **14**, 719-734, 1985.
- Law, A.M., and W.D. Kelton, *Simulation Modeling and Analysis*, McGraw-Hill, 3rd edition, 2000.
- Law, A.M., and M.G. McComas, Simulation-Based Optimization, *Proc. Winter Simulation Conference, San Diego, CA*, 41-44, 2002.
- Nissen, V., and J. Propach, On the Robustness of Population-Based Versus Point-Based Optimization in the Presence of Noise, *IEEE Transactions on Evolutionary Computation* **2**, 107-119, 1998.
- Rinott, Y., On Two-Stage Selection Procedures and Related Probability-Inequalities, *Communications in Statistics – Theory and Methods* **A7**, 799-811, 1978.
- Schwefel, H.P., *Evolution and Optimum Seeking*, John Wiley, 1995.
- Striver, T.A., and J.W. Chrissis, Combined Pattern Search and Ranking and Selection for Simulation Optimization, *Proc. Winter Simulation Conference, Washington, DC*, 645-653, 2004.
- Swisher, J.R., S.H. Jacobson, and E. Yücesan, Discrete-Event Simulation Optimization Using Ranking, Selection, and Multiple Comparison Procedures: A Survey, *ACM Transactions on Modeling and Computer Simulation* **13**, 134-154, 2003.

AUTHOR BIOGRAPHIES

PETER BUCHHOLZ is a professor for Modeling and Simulation at the University of Dortmund. Previously he held a professorship at Dresden University of Technology. He holds a diploma, Dr. rer. nat. (Ph.D.) and a habilitation degree, all from the University of Dortmund. His current research interests are analysis techniques for discrete event systems including simulation, the combination of analysis and optimisation and the application of the methods to communication systems and logistic networks. His Web address is <www4.cs.uni-dortmund.de/~Buchholz>.

AXEL THÜMMLER is a research scientist in the Modeling and Simulation group at the University of Dortmund. He received the degree Diplom-Informatiker (M.S. in Computer Science) in April 1998 and the degree Dr. rer. nat. (Ph.D. in Computer Science) in July 2003, both from the University of Dortmund. His research interests include simulation optimization, communication networks, mobile computing, and performance evaluation. His Web address is <www4.cs.uni-dortmund.de/~Thuemmler>.