Optimum Tracking with Evolution Strategies

Dirk V. Arnold and Hans-Georg Beyer
Department of Computer Science XI
University of Dortmund
44221 Dortmund, Germany
{arnold,beyer}@Ls11.cs.uni-dortmund.de

Abstract

Evolutionary algorithms are frequently applied to dynamic optimization problems in which the objective varies with time. It is desirable to gain an improved understanding of the influence of different genetic operators and of the parameters of a strategy on its tracking performance. An approach that has proven useful in the past is to mathematically analyze the strategy's behavior in simple, idealized environments. The present paper investigates the performance of a multiparent evolution strategy that employs cumulative step length adaptation for an optimization task in which the target moves linearly with uniform speed. Scaling laws that quite accurately describe the behavior of the strategy and that greatly contribute to its understanding are derived. It is shown that in contrast to previously obtained results for a randomly moving target, cumulative step length adaptation fails to achieve optimal step lengths if the target moves in a linear fashion. Implications for the choice of population size parameters are discussed.

Keywords: genetic and evolutionary computation, step length adaptation, tracking problem, dynamic optimization

1 Introduction

Evolutionary algorithms are nature inspired heuristics for search and optimization that model the iterated interplay of variation and selection in a population of candidate solutions. Comprehensive treatments of the subject area can be found in (Bäck, 1996; Goldberg, 1989; Mitchell, 1996; Rechenberg, 1994). Reasons for the widespread use of evolutionary algorithms include their often observed robustness, the ease with which the underlying paradigm is understood and implemented, and their wide applicability. Areas of application today include management, control, design, scheduling, pattern recognition, and decision making.

In recent years, a fair amount of theoretical investigation has contributed substantially to our understanding of the dynamics of evolutionary search strategies on a variety of problem classes. For an overview, see (Beyer et al., 2002). However, most of the problem classes that have been considered are of a static nature. In contrast, many problems encountered in the computational, engineering, and biological sciences are dynamic in that the objective is not constant but varies with time. Instances of dynamic optimization problems arise for example in the context of online job scheduling, where new jobs arrive in the course of the optimization. In the engineering sciences, many control problems are of an inherently dynamic nature. An extensive list of references concerned with dynamic optimization problems in the biological sciences has been reviewed by (Bürger, 2000).

In contrast to static optimization, where the goal is to rapidly and accurately locate a (near) optimal solution, the task in dynamic optimization frequently amounts to tracking a moving target as closely as possible. Strategies for dynamic optimization need to continually adapt to changes in the environment. Evolutionary algorithms are strategies that are population based rather than point based, and that place great emphasis on adaptability. They are thus believed to be good at handling time-varying objectives. In fact, dynamic optimization is frequently cited as a prime application area for evolutionary algorithms.

The knowledge available with respect to the capabilities of evolutionary algorithms for dynamic problems is mostly of an empirical nature. An extensive survey of work concerned with evolutionary optimization in dynamic environments along with a collection of benchmark functions and a discussion of methods that have been proposed to improve the performance of evolutionary algorithms when the objective varies with time has been compiled by (Branke, 2001). (Angeline, 1997) compares empirically the tracking performance of an evolutionary algorithm employing a form of mutative self-adaptation with that of a strategy using a simple heuristic for step length adaptation. The fitness environment considered is a three-dimensional, spherically symmetric objective function that is shifted periodically either in a random fashion or on a linear or a spherical path. Angeline observes that the self-adaptation mechanism is not without problems in the dynamic case. In that same fitness environment, (Bäck, 1998) compares different variants of mutative self-adaptation and presents evidence that seems to indicate that the lognormal self-adaptation used in evolution strategies performs better than the variant of self-adaptation commonly used in evolutionary programming. (Salomon and Eggenberger, 1997) compare empirically the performance of evolution strategies with that of a breeder genetic algorithm on the sphere, an ellipsoid, and Rastrigin's function, where the coordinates are shifted by a constant increment in every time step. The search space dimensionalities they consider for the sphere are $N = 10$ and $N = 30$. While they do not quantify the term, they find that the sensitivity to the particular implementation of the strategy and to its parameter values is much lower for the tracking task than it is in a static environment. Without providing details, they also report to have observed that recombination is not beneficial for tracking a moving target. (Weicker and Weicker, 1999) contrast self-adaptation of a single mutation strength with that of N mutation strengths and adaptation of the full mutation covariance matrix and find that in more rapidly changing environments, the adaptation of more than a single mutation strength becomes unreliable. Finally, (Droste, 2002) presents a first rigorous analysis of the performance of a $(1+1)$-strategy on a discrete, dynamic objective function. However, focus in that paper is not on the tracking behavior of the strategy but rather on the expected time required to first reach the optimum.

Altogether, the above references provide a host of empirical knowledge with respect to evolutionary optimization in dynamic environments. In an attempt to establish some theoretical results to complement and explain that empirical knowledge, in (Arnold and Beyer, 2002b) the performance of a multiparent evolution strategy with cumulative step length adaptation has been analyzed for a dynamic optimization problem in which the target moves in a random fashion. The tools employed in the work had previously been used in an analysis of the behavior of that strategy on a static optimization problem disturbed by noise (Arnold, 2002; Arnold and Beyer, 2002a). It was found that for the case of the randomly moving target, cumulative step length adaptation is able to achieve asymptotically optimal step lengths. It was noted that the case of a target that moves linearly rather than randomly is an interesting one to consider as the linear motion introduces positive correlations in the sequence of steps to be taken. As cumulative step length adaptation relies on the conjecture that consecutive steps should be uncorrelated, it appears reasonable to expect suboptimal performance of cumulative step length adaptation for the case of the linearly moving target.

The present article presents the analysis of the tracking performance of the $(\mu/\mu, \lambda)$-ES with cumulative step length adaptation for the case of a linearly moving target. The analysis is more difficult than that of the tracking problem with random motion of the target as the linear motion of the target introduces an additional degree of freedom in the calculations. The fitness environment closely resembles that considered in (Angeline, 1997; Bäck, 1998; Salomon and Eggenberger, 1997). However, in contrast to those references, no particular values need to be assumed for the size of the population, the speed of the target, or the search space dimensionality. Instead, scaling laws that describe the dependence of the performance of the strategy on those parameters are obtained. Such scaling laws provide a quantitative understanding of the dynamics of the evolutionary processes that makes it possible to predict optimal parameter settings for the problem class considered. Moreover, we share the belief of (van Nimwegen and Crutchfield, 2001) that a general, predictive theory of the dynamics of evolutionary search can be built incrementally, starting with a quantitative analytical understanding of specific problems and then generalizing to more complex situations. In this view, the present article can be considered as a step neither at the beginning nor at the end of the chain of problems to be considered.

2 The $(\mu/\mu, \lambda)$-ES

In contrast to genetic algorithms that are most often used in discrete search spaces, evolution strategies are a class of evolutionary algorithms that are used mostly for the optimization of real-valued functions $f : I\!R^N \to I\!R$. The $(\mu/\mu, \lambda)$-ES is a variant that is popular both due to its proven good performance in static settings and its relative mathematical tractability. For a thorough introduction to evolution strategies, see (Bäck, 1996; Rechenberg, 1994; Schwefel, 1995). An explanation of the $(\mu/\mu, \lambda)$-symbolism can be found in (Beyer, 2001).

Like other evolutionary algorithms, the $(\mu/\mu, \lambda)$-ES strives to drive a population of candidate solutions towards increasingly better regions of the search space. It operates with a population $\mathcal{P}$ of μ candidate solutions. In every time step t, a set $\mathcal{Q}^{(t)}$ of λ candidate solutions is generated from $\mathcal{P}^{(t)}$ by means of variation. Subsequently, selection ensures that the μ best of the λ candidate solutions in $\mathcal{Q}^{(t)}$ form $\mathcal{P}^{(t+1)}$. Figure 1 illustrates the basic procedure. Note that while generally, initialization schemes and termination criteria are important components of the algorithm, they are frequently application dependent and are irrelevant in the present context. Rather than considering them here, we refer to (Bäck, 1996) for a discussion.

Variation is accomplished by means of recombination and mutation. For the $(\mu/\mu, \lambda)$-ES, recombination simply computes the centroid, i.e. the arithmetic mean, of the parental population. Mutation consists of adding a random vector $\sigma\mathbf{z}$, where $\mathbf{z}$ consists of independent, standard normally distributed components, to that centroid.[1] The standard deviation σ of the components of vector $\sigma\mathbf{z}$ is referred to as the mutation strength, vector $\mathbf{z}$ as a mutation vector. Notice that the mutation strength determines the step length of the strategy. The average of those mutation vectors that correspond to offspring candidate solutions that are selected to form the population of the next time step is the progress vector $\langle\mathbf{z}\rangle$. Letting $\mathbf{x}$ denote the population centroid, selection thus ensures that

$$\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)} + \sigma\langle\mathbf{z}\rangle^{(t)} \tag{1}$$

holds.

[1]The restriction to isotropic mutations has been made here in order to keep the analysis of the strategy's behavior tractable. Practical implementations usually rely on correlated mutations that can vary in magnitude in different dimensions. The step length adaptation scheme outlined below has a generalization for that case that is described in (Hansen, 1998; Hansen and Ostermeier, 2001).

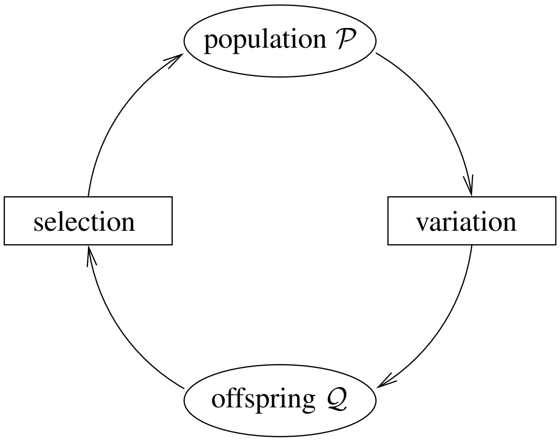

Figure 1: *The basic evolution loop. Operators that introduce variation are applied to a population $\mathcal{P}$ of μ candidate solutions to generate a set $\mathcal{Q}$ of $\lambda > \mu$ new candidate solutions. The μ best of the newly generated candidate solutions are then selected to form the population of the next time step.*

In real-valued search spaces, it is generally necessary for the mutation strength to be adapted continuously to the local characteristics of the objective function. Information from the present and possibly from past time steps needs to be used in order to learn appropriate step sizes. A mechanism that is commonly employed for the adaptation of the mutation strength is the cumulative step length adaptation algorithm of (Hansen, 1998; Hansen and Ostermeier, 2001). It relies on the conjecture that if the mutation strength is below its optimal value, consecutive steps of the strategy tend to be parallel, and if the mutation strength is too high, consecutive steps tend to be antiparallel. For optimally adapted mutation strength, the steps taken by the evolution strategy are uncorrelated. This is plausible intuitively as several steps in the same direction in search space are ideally replaced by a single longer step in that direction. Conversely, consecutive steps that nullify each other are a sign that the step length is too high. So as to be able to reliably detect parallel or antiparallel correlations between successive steps, information from a number of time steps needs to be accumulated. For the $(\mu/\mu, \lambda)$-ES, the accumulated progress vector $\mathbf{s}$ is defined by $\mathbf{s}^{(0)} = \mathbf{0}$ and the recursive relationship

$$\mathbf{s}^{(t+1)} = (1-c)\mathbf{s}^{(t)} + \sqrt{\mu c(2-c)}\langle \mathbf{z} \rangle^{(t)}, \tag{2}$$

where c is a constant determining how far back the "memory" of the accumulation process reaches. The mutation strength is updated according to

$$\sigma^{(t+1)} = \sigma^{(t)} \exp\left(\frac{\|\mathbf{s}^{(t+1)}\|^2 - N}{2DN} \right), \tag{3}$$

where D denotes a damping constant. Note that the term N in the numerator of the argument to the exponential function is the mean squared length of the accumulated progress vector if consecutive progress vectors are stochastically independent. If the squared length of the accumulated progress vector is less than N then the mutation strength is decreased. If it is greater than N then the mutation strength is increased. Also note that the prescription Eq. (3) for adapting the mutation strength has been changed slightly from the prescription in the original algorithm given by (Hansen, 1998) in that here, adaptation is performed on the basis of the squared length of the accumulated progress vector rather than on its length. The difference in performance appears to be insignificant while elegance in the formulation is gained by the change. The constants c and D are set to $1/\sqrt{N}$ and $\sqrt{N}$, respectively, according to recommendations made by (Hansen, 1998).

It should be mentioned that cumulative step length adaptation is not the only mechanism conceivable for the adaptation of mutation strengths. Possible alternatives include:

Nested evolution strategies (Herdy, 1992; Rechenberg, 1994). Nested evolution strategies adjust strategy parameters such as mutation strengths by means of evolutionary optimization on a meta level. Several populations, each one with their own parameter settings, compete with each other for survival. After a number of time steps, the respective progress of the different strategies is examined. The mutation strengths of those populations that have achieved the largest progress are used as a basis for generating mutation strengths for the next round of competition by means of recombination and mutation.

Mutative self adaptation (Rechenberg, 1994; Schwefel, 1995). Mutative self-adaptation includes the mutation strengths into the optimization process at the same hierarchical level as the object parameters of the problem. Different candidate solutions have differing mutation strengths. Assuming that favorable mutation strengths are more likely to generate successful offspring than unfavorable ones, selection of favorable mutation strengths is then a by-product of evolution.

Machine learning approaches (Ravisé and Sebag, 1996; Sebag et al., 1997). The problem of mutation strength adaptation can be understood as a learning task. An explicit memory of the evolutionary process can be created in the form of a set of rules by discriminating between successful trials and unsuccessful ones. Mutation strengths can then be adapted by means of inductive learning. (Michalski, 2000) suggests a different approach in which a machine learning system seeks reasons why certain individuals in a population are superior to others. These reasons, formulated as inductive hypotheses, are then used to generate new populations directly rather than by recombination and mutation.

In the present article, we choose to consider cumulative step length adaptation due to its relative mathematical tractability and as it is known to generate asymptotically optimal step lengths for the random tracking problem studied in (Arnold and Beyer, 2002b).

3 The Tracking Problem

The purpose of this section is to introduce the tracking problem studied in the remainder of this article. In order to study the tracking performance of the $(\mu/\mu, \lambda)$-ES, we consider a dynamic version of the sphere model. The sphere model is the set of all functions $f : \mathbb{R}^N \to \mathbb{R}$ with

$$f(\mathbf{x}) = g(\|\hat{\mathbf{x}} - \mathbf{x}\|),$$

where $g : \mathbb{R} \to \mathbb{R}$ is a strictly monotonic function of the distance $R = \|\mathbf{R}\| = \|\hat{\mathbf{x}} - \mathbf{x}\|$ of a candidate solution $\mathbf{x}$ from the target $\hat{\mathbf{x}}$. It has frequently served as a model for fitness landscapes at a stage where the population of candidate solutions is in relatively close proximity to the target and is most often studied in the limit of very high search space dimensionality. So as to study the tracking behavior of evolutionary algorithms, several authors (Angeline, 1997; Arnold and Beyer, 2002b; Bäck, 1998; Salomon and Eggenberger, 1997) have added a dynamic component to the sphere model by stipulating that the target $\hat{\mathbf{x}}$ vary with time. Several modes of motion of the target are conceivable and have been explored. Examples include random motion, linear motion, and circular motion in search space. The

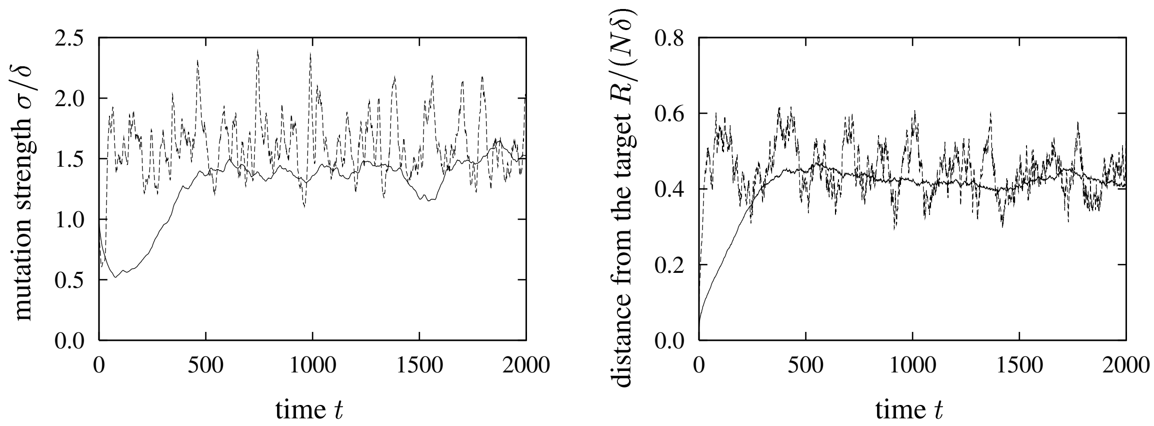

Figure 2: *Mutation strength σ/δ and stationary distance from the target $R/(N\delta)$ as functions of the time t. The strategy is a $(3/3, 10)$-ES with cumulative step length adaptation. The population is initialized to be centered at the location of the target at time step 0. The initial mutation strength is $\sigma = \delta$. The dashed and solid lines have been measured in typical runs of the strategy with search space dimensionalities $N = 40$ and $N = 400$, respectively.*

case of random motion has been studied from a theoretical perspective in (Arnold and Beyer, 2002b). If the target shifts such that

$$\hat{\mathbf{x}}^{(t+1)} = \hat{\mathbf{x}}^{(t)} + \delta\hat{\mathbf{z}},$$

where vector $\hat{\mathbf{z}}$ consists of N independent, standard normally distributed components, then it was seen that the mutation strength that minimizes the distance at which the target is tracked is $\sigma = \sqrt{\mu}\delta$, and that in the limit $N \to \infty$, that mutation strength is realized by cumulative step length adaptation. In the present paper, we consider linear motion and assume that the target at time step $t + 1$ is

$$\hat{\mathbf{x}}^{(t+1)} = \hat{\mathbf{x}}^{(t)} + \delta\mathbf{v}, \tag{4}$$

where vector $\mathbf{v}$ is constant and has unit length. As in the case of random motion, the factor δ is a measure for the speed of change of the objective.

Figure 2 illustrates the time behavior of a $(\mu/\mu, \lambda)$-ES with $\mu = 3$ and $\lambda = 10$ that uses cumulative step length adaptation for the tracking problem with linear dynamics of the target. Shown are measurements of typical runs of the strategy for search space dimensionalities $N = 40$ and $N = 400$. The results that are obtained for other values of μ and λ qualitatively agree. It can be seen that after a transient period, a state that is characterized by the mutation strength and the distance from the target fluctuating around stationary mean values is reached. Additional measurements not included here seem to indicate that the length of the transient period is of order N. The stationary state that is reached is independent of the initialization of the strategy. Furthermore, it can be seen that the magnitude of the fluctuations decreases with increasing search space dimensionality. It is the goal in the remainder of this paper to characterize the dependence of the stationary state on the population size parameters μ and λ, the speed of the target δ, and the search space dimensionality N.

Analyses of the behavior of evolution strategies on the sphere model rely on a decomposition of vectors that is illustrated in Figure 3. A vector $\mathbf{z}$ originating at search space location $\mathbf{x}$ can be written as the sum of two vectors $\mathbf{z}_A$ and $\mathbf{z}_B$, where $\mathbf{z}_A$ is parallel to $\mathbf{R} = \hat{\mathbf{x}} - \mathbf{x}$ and $\mathbf{z}_B$ is in the hyperplane perpendicular to that. In the present context, $\mathbf{z}$ can be either a mutation vector or a progress vector. The vectors $\mathbf{z}_A$ and $\mathbf{z}_B$ are referred to as the central and lateral components of vector $\mathbf{z}$, respectively.

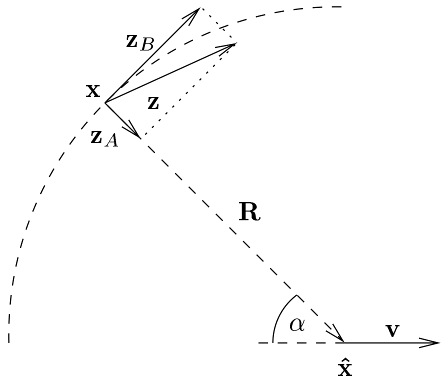

Figure 3: *Decomposition of a vector $\mathbf{z}$ into central component $\mathbf{z}_A$ and lateral component $\mathbf{z}_B$. Vector $\mathbf{z}_A$ is parallel to $\hat{\mathbf{x}} - \mathbf{x}$, vector $\mathbf{z}_B$ is in the hyperplane perpendicular to that. The trailing angle α is the angle formed by vector $\mathbf{R} = \hat{\mathbf{x}} - \mathbf{x}$ and the direction $\mathbf{v}$ of the motion of the target.*

The signed length z_A of the central component of vector $\mathbf{z}$ is defined to equal $\|\mathbf{z}_A\|$ if $\mathbf{z}_A$ points towards the target and to equal $-\|\mathbf{z}_A\|$ if it points away from it.

Due to the stochastic nature of the search process and in particular to the presence of the lateral component of mutation vectors, the population follows the target not on the straight line defined by vector $\mathbf{v}$ but at an angle that is nonzero on average. As illustrated in Figure 3, the trailing angle α is defined as the angle formed by the vector $\mathbf{R} = \hat{\mathbf{x}} - \mathbf{x}$, where $\mathbf{x}$ is the population centroid, and the direction $\mathbf{v}$ of the motion of the target. The analyses in Sections 4 and 5 will reveal that the average magnitude of that angle can be substantial.

Clearly, from Eqs. (1) and (4),

$$\mathbf{R}^{(t+1)} = \mathbf{R}^{(t)} + \delta\mathbf{v} - \sigma\langle\mathbf{z}\rangle^{(t)} \tag{5}$$

holds as a result of the motion of the target and the step that the strategy takes. Moreover, it is well known from (Beyer, 2001; Rechenberg, 1994) that in the limit $N \to \infty$, the expectation of the signed length of the central component of the progress vector is

$$\mathrm{E}\left[\langle z_A\rangle\right] = c_{\mu/\mu,\lambda}, \tag{6}$$

where $c_{\mu/\mu,\lambda}$ is the expectation of the average of the μ last order statistics of a sample of λ independent, standard normally distributed random variables. An explicit prescription for computing $c_{\mu/\mu,\lambda}$ has been provided in (Beyer, 2001). The lateral component of the progress vector is of random direction in the hyperplane defined by normal vector $\mathbf{R} = \hat{\mathbf{x}} - \mathbf{x}$. Moreover, in the limit of infinite search space dimensionality, the expected overall squared length of the progress vector is

$$\mathrm{E}\left[\|\langle\mathbf{z}\rangle\|^2\right] = \frac{N}{\mu}. \tag{7}$$

The standard deviation of that squared length is of order $\sqrt{N}$, diminishing the relative influence of fluctuations as N tends to infinity. Notice that this effect is responsible for the decrease in fluctuations that had been observed in Figure 2. From Eq. (6) and Figure 3 it is easily seen that

$$\mathrm{E}[\mathbf{R}\cdot\langle\mathbf{z}\rangle] = Rc_{\mu/\mu,\lambda} \tag{8}$$

holds. Moreover, due to the randomness of the direction of the lateral component of the progress vector,

$$\begin{aligned}\mathrm{E}[\langle\mathbf{z}\rangle\cdot\mathbf{v}] &= \mathrm{E}[\langle\mathbf{z}_A\rangle\cdot\mathbf{v}] + \underbrace{\mathrm{E}[\langle\mathbf{z}_B\rangle\cdot\mathbf{v}]}_{=0} \\ &= c_{\mu/\mu,\lambda}\cos\alpha\end{aligned} \tag{9}$$

holds as a consequence of Eq. (6).

4 Analysis for Constant Mutation Strength

Before proceeding to the analysis of the behavior of the $(\mu/\mu,\lambda)$-ES with cumulative step length adaptation, in this section we first consider the case that the mutation strength is fixed. The strategy in combination with the dynamic objective function outlined in Section 3 forms a stochastic dynamic system. Due to the symmetries inherent in both the environment and the strategy, for fixed mutation strength, the state of that system can described by just two variables. A possible choice for those variables are the distance R of the population centroid from the target and the cosine of the trailing angle α. Provided that the mutation strength is large enough for the strategy to be able to keep pace with the target, the system tends toward a limit state in which the state variables have time-invariant distributions.

The approach to determining the scaling behavior of the $(\mu/\mu,\lambda)$-ES on the tracking problem with linear dynamics of the target is the same as that employed in (Arnold and Beyer, 2002b) for the tracking problem with random dynamics of the target. It consists in using stationarity conditions while neglecting all fluctuations. In particular, it is assumed that the progress vector has a central component of length $c_{\mu/\mu,\lambda}$, that its lateral component is of random direction, and that its overall squared length is N/μ. As a consequence, quantities involving the progress vector can be replaced by their expectations given in Eqs. (6), (7), (8), and (9). The state variables tend to stationary limit values that can be obtained by relatively simple considerations. As fluctuations of the progress vector decrease with increasing search space dimensionality, the approach that neglects them becomes increasingly accurate with increasing N. We will see that good predictions are obtained already for relatively moderate values of N.

It can be seen from Figure 3 that $\mathbf{R}\cdot\mathbf{v} = R\cos\alpha$. Therefore, at time step $t+1$,

$$\begin{aligned}R^{(t+1)}\cos\alpha^{(t+1)} &= \mathbf{R}^{(t+1)}\cdot\mathbf{v} \\ &= \left(\mathbf{R}^{(t)} + \delta\mathbf{v} - \sigma\langle\mathbf{z}\rangle^{(t)}\right)\cdot\mathbf{v} \\ &\simeq R^{(t)}\cos\alpha^{(t)} + \delta - \sigma c_{\mu/\mu,\lambda}\cos\alpha^{(t)},\end{aligned}$$

where Eqs. (5) and (9) have been used. As stationarity of the state variables implies the stationarity of their product, it follows

$$\cos\alpha \simeq \frac{\delta}{\sigma c_{\mu/\mu,\lambda}} \tag{10}$$

for the stationary cosine of the trailing angle.

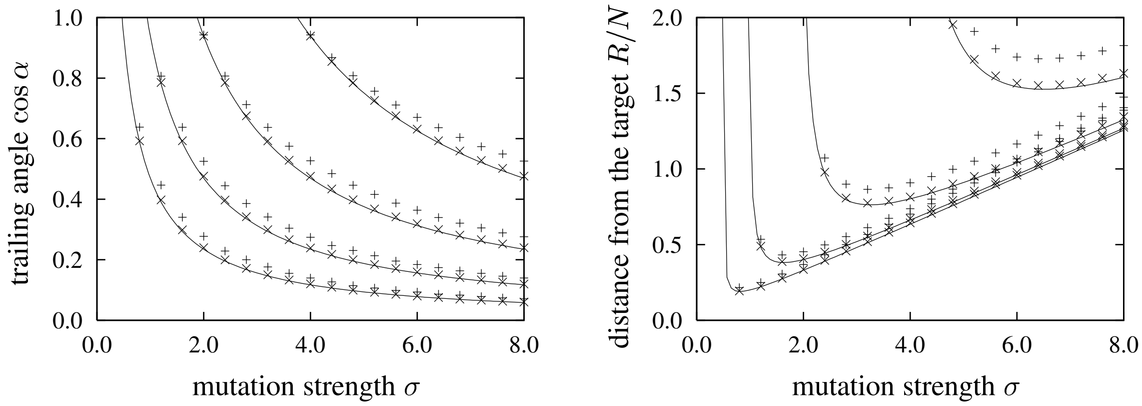

Figure 4: *Cosine of the trailing angle and distance from the target as functions of the mutation strength* σ. *The target moves linearly with, from bottom to top, speeds* $\delta = 0.5$, 1.0, 2.0, *and* 4.0. *The strategy is a* $(3/3, 10)$*-ES. The solid lines are predictions from Eqs.* (10) *and* (11), *the points mark empirical measurements for search space dimensionalities* $N = 40$ $(+)$ *and* $N = 400$ $(\times)$.

Similarly, the square of the distance from the population centroid to the target at time step $t+1$ is

$$\begin{aligned} {R^{(t+1)}}^2 &= \mathbf{R}^{(t+1)} \cdot \mathbf{R}^{(t+1)} \\ &= \left(\mathbf{R}^{(t)} + \delta\mathbf{v} - \sigma\langle\mathbf{z}\rangle^{(t)}\right) \cdot \left(\mathbf{R}^{(t)} + \delta\mathbf{v} - \sigma\langle\mathbf{z}\rangle^{(t)}\right) \\ &\simeq {R^{(t)}}^2 + \delta^2 + \frac{N\sigma^2}{\mu} + 2\delta R^{(t)} \cos\alpha^{(t)} - 2\sigma R^{(t)} c_{\mu/\mu,\lambda} - 2\delta\sigma c_{\mu/\mu,\lambda} \cos\alpha^{(t)}, \end{aligned}$$

where Eqs. (5), (7), (8), (9), and the fact that $\mathbf{R} \cdot \mathbf{v} = R\cos\alpha$ have been used. From the stationarity of R it follows that R^2 is stationary. Using Eq. (10), and taking into account that the influence of the second and sixth terms in the last line vanishes in the limit $N \to \infty$ yields[2]

$$R \simeq \frac{N\sigma^3 c_{\mu/\mu,\lambda}}{2\mu(\sigma^2 c_{\mu/\mu,\lambda}^2 - \delta^2)} \tag{11}$$

for the stationary distance from the population centroid to the target.

The quality of the approximation to the stationary state thus derived is illustrated in Figure 4. Predictions from Eqs. (10) and (11) are compared with empirical measurements of runs of evolution strategies. For those measurements, it is irrelevant whether results from one run are used or several runs are averaged, provided that the length of the runs is sufficient. For the experiments reported here, the length of the runs was chosen to ensure that the standard deviation of the measurements is below the size of the crosses in Figure 4. It can be seen from the figure that for search space dimensionality $N = 40$, the qualitative dependence of the performance of the strategy on the mutation strength and the target speed is reflected properly. The deviations that can be observed are attributable to

[2] In the previous equation, assuming invariance, i.e. $R^{(t+1)} = R^{(t)} = R$, using Eq. (10) to replace $\cos\alpha$, and rearranging terms yields

$$2R\left(\sigma c_{\mu/\mu,\lambda} - \frac{\delta^2}{\sigma c_{\mu/\mu,\lambda}}\right) \simeq \frac{N\sigma^2}{\mu}\left(1 - \frac{\mu\delta^2}{N\sigma^2}\right).$$

In order for the strategy to be able to track the target, the mutation strength must be large enough to satisfy $\delta/\sigma = \mathcal{O}(1)$. Therefore, the last term on the right hand side vanishes as $N \to \infty$ and Eq. (11) immediately follows by solving for R.

the fluctuations inherent in the runs that have been neglected in the calculations. For search space dimensionality $N = 400$, very good quantitative agreement of the predictions with the empirical measurements can be observed.

Some important conclusions with regard to the scaling behavior of the strategy can be drawn from Eqs. (10) and (11). First of all, from Eq. (11) it can be seen that a stationary state is reached only if $\sigma > \delta/c_{\mu/\mu,\lambda}$. For smaller mutation strengths, the strategy is not able to keep pace with the target and the distance from the population centroid to the target increases indefinitely. Second, Eq. (11) can be used to obtain the optimal mutation strength, i.e. the mutation strength that minimizes the distance to the target. Computing the derivative with respect to σ and determining the root yields optimal mutation strength

$$\sigma \simeq \frac{\sqrt{3}\delta}{c_{\mu/\mu,\lambda}}. \tag{12}$$

Using Eq. (11), the resulting distance from the target is

$$R \simeq \frac{3\sqrt{3}N\delta}{4\mu c_{\mu/\mu,\lambda}^2}. \tag{13}$$

Therefore, combining Eqs. (12) and (13), we have

$$\sigma\frac{N}{R} \simeq \frac{4}{3}\mu c_{\mu/\mu,\lambda} \tag{14}$$

for the optimal mutation strength. Comparison with the corresponding result from (Arnold and Beyer, 2002b) reveals that this differs by a factor of $4/3$ from the optimal mutation strength for a randomly moving target. As in the case of a linearly moving target the cosine of the trailing angle is consistently greater than zero, the strategy should ideally operate with a larger mutation strength.

5 Analysis for Adaptive Mutation Strength

For a stationary or a randomly moving target, cumulative step length adaptation has been seen to be able to generate asymptotically optimal mutation strengths (Arnold, 2002; Arnold and Beyer, 2002a; Arnold and Beyer, 2002b). In the present section, we study the performance of cumulative step length adaptation for the case of the linearly moving target.

For fixed mutation strength, we have seen in the previous section that a system state is fully characterized by two state variables. When using cumulative step length adaptation, a total of three state variables beyond those introduced in Section 4 are required to describe the state of the system. While again several choices are conceivable, one possibility is to consider the mutation strength σ along with the signed length s_A of the central component of the accumulated progress vector and that vector's overall squared length $\|\mathbf{s}\|^2$. Note that the same choice has been made in (Arnold, 2002; Arnold and Beyer, 2002a; Arnold and Beyer, 2002b). In analogy to Eq.(9),

$$\begin{aligned} \mathrm{E}[\mathbf{s}\cdot\langle\mathbf{z}\rangle] &= \mathrm{E}[s_A\langle z_A\rangle] + \underbrace{\mathrm{E}[\mathbf{s}_B\cdot\langle\mathbf{z}_B\rangle]}_{=0} \\ &= s_A c_{\mu/\mu,\lambda} \end{aligned} \tag{15}$$

holds due to Eq. (6) and the randomness of the direction of the lateral component of progress vectors.

As in Section 4, we neglect fluctuations of the progress vector and make use of the stationarity of the limit values of the state variables. The inner product $\mathbf{s}\cdot\mathbf{v}$ at time step $t+1$ is

$$\begin{aligned}\mathbf{s}^{(t+1)}\cdot\mathbf{v} &= \Big((1-c)\mathbf{s}^{(t)} + \sqrt{\mu c(2-c)}\langle\mathbf{z}\rangle^{(t)}\Big)\cdot\mathbf{v}\\ &\simeq (1-c)\mathbf{s}^{(t)}\cdot\mathbf{v} + \sqrt{\mu c(2-c)}c_{\mu/\mu,\lambda}\cos\alpha^{(t)},\end{aligned}$$

where Eqs. (2) and (9) have been used. From the stationarity of the state variables follows the stationarity of $\mathbf{s}\cdot\mathbf{v}$ and therefore with Eq. (10)

$$\mathbf{s}\cdot\mathbf{v} \simeq \sqrt{\frac{\mu(2-c)}{c}}\frac{\delta}{\sigma} \tag{16}$$

for the stationary value of the inner product of vectors $\mathbf{s}$ and $\mathbf{v}$.

Clearly, due to the definition of the central component of vectors, $\mathbf{R}\cdot\mathbf{s} = Rs_A$ holds. Therefore, at time step $t+1$,

$$\begin{aligned}R^{(t+1)}s_A^{(t+1)} &= \mathbf{R}^{(t+1)}\cdot\mathbf{s}^{(t+1)}\\ &= \Big(\mathbf{R}^{(t)} + \delta\mathbf{v} - \sigma^{(t)}\langle\mathbf{z}\rangle^{(t)}\Big)\cdot\Big((1-c)\mathbf{s}^{(t)} + \sqrt{\mu c(2-c)}\langle\mathbf{z}\rangle^{(t)}\Big)\\ &\simeq (1-c)R^{(t)}s_A^{(t)} + (1-c)\sqrt{\frac{\mu(2-c)}{c}}\frac{\delta^2}{\sigma^{(t)}} - (1-c)\sigma^{(t)}c_{\mu/\mu,\lambda}s_A^{(t)}\\ &\qquad + \sqrt{\mu c(2-c)}R^{(t)}c_{\mu/\mu,\lambda} + \sqrt{\mu c(2-c)}\frac{\delta^2}{\sigma^{(t)}} - \sqrt{\mu c(2-c)}\frac{N\sigma^{(t)}}{\mu},\end{aligned}$$

where Eqs. (2), (5), (7), (8), (9), (10), (15), and (16) have been used. Of the six terms in the last line, the influence of the second, third, and fifth vanishes in the limit of infinite search space dimensionality.[3] Therefore, from the remaining terms, demanding stationarity and using Eq. (11) yields

$$\begin{aligned}s_A &\simeq \sqrt{\frac{\mu(2-c)}{c}}\left(c_{\mu/\mu,\lambda} - \frac{N\sigma}{R\mu}\right)\\ &\simeq \sqrt{\frac{\mu(2-c)}{c}}\left(\frac{2\delta^2}{c_{\mu/\mu,\lambda}\sigma^2} - c_{\mu/\mu,\lambda}\right)\end{aligned} \tag{17}$$

for the stationary value of the signed length of the central component of the accumulated progress vector.

Similarly, the squared length of the accumulated progress vector at time step $t+1$ is

$$\begin{aligned}\|\mathbf{s}^{(t+1)}\|^2 &= \Big((1-c)\mathbf{s}^{(t)} + \sqrt{\mu c(2-c)}\langle\mathbf{z}\rangle^{(t)}\Big)\cdot\Big((1-c)\mathbf{s}^{(t)} + \sqrt{\mu c(2-c)}\langle\mathbf{z}\rangle^{(t)}\Big)\\ &= (1-c)^2\|\mathbf{s}^{(t)}\|^2 + 2(1-c)\sqrt{\mu c(2-c)}\mathbf{s}^{(t)}\cdot\langle\mathbf{z}\rangle^{(t)} + \mu c(2-c)\|\langle\mathbf{z}\rangle^{(t)}\|^2\\ &\simeq (1-c)^2\|\mathbf{s}^{(t)}\|^2 + 2(1-c)\sqrt{\mu c(2-c)}c_{\mu/\mu,\lambda}s_A^{(t)} + c(2-c)N,\end{aligned}$$

[3] Assuming stationarity, i.e. $R^{(t+1)}s_A^{(t+1)} = R^{(t)}s_A^{(t)} = Rs_A$, and rearranging terms yields

$$s_A\left(c + (1-c)\frac{\sigma c_{\mu/\mu,\lambda}}{R}\right) \simeq \sqrt{\mu c(2-c)}\left(\frac{\delta^2}{cR\sigma} + c_{\mu/\mu,\lambda} - \frac{N\sigma}{R\mu}\right).$$

As seen in Section 4, in order for the strategy to be able to track the target, $\delta/\sigma = \mathcal{O}(1)$ and $\sigma N/R = \mathcal{O}(1)$ must hold. Therefore, the second term in the parentheses on the left hand side and the first term in the parentheses on the right hand side vanish as $N\to\infty$. Rearranging the remaining terms immediately yields Eq. (17)

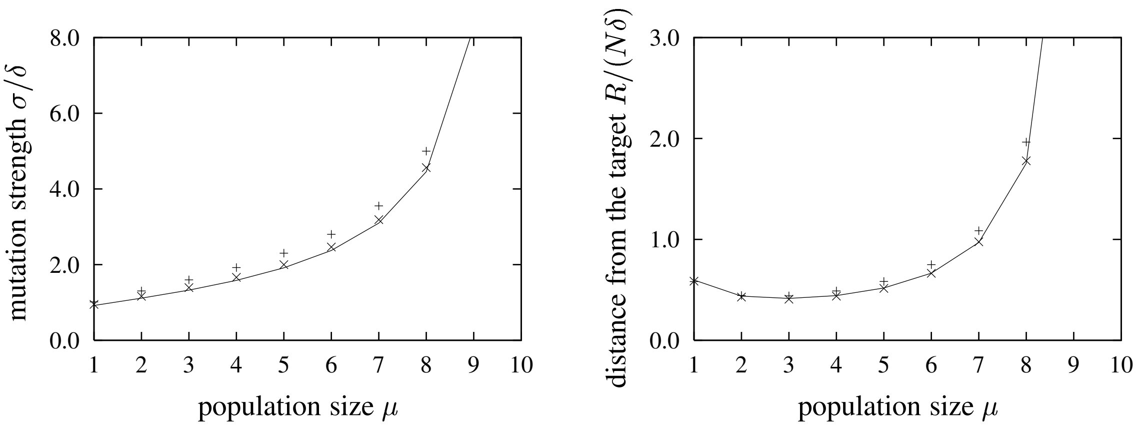

Figure 5: *Mutation strength σ/δ and stationary distance from the target $R/(N\delta)$ as functions of the population size μ. The strategy is a $(\mu/\mu, 10)$-ES. The solid lines are obtained from Eqs.* (20) *and* (21). *The points mark empirical measurements for search space dimensionalities $N = 40$ $(+)$ and $N = 400$ $(\times)$.*

where Eqs. (2), (7), and (15) have been used. Demanding stationarity and using Eq. (17) yields

$$\|\mathbf{s}\|^2 \simeq 2\frac{1-c}{c}c_{\mu/\mu,\lambda}\left(\mu c_{\mu/\mu,\lambda} - \frac{\sigma N}{R}\right) + N \tag{18}$$

for the stationary squared length of the accumulated progress vector.

Finally, demanding stationarity of the mutation strength and using Eqs. (3) and (18) yields

$$\sigma\frac{N}{R} \simeq \mu c_{\mu/\mu,\lambda} \tag{19}$$

for the mutation strength that cumulative mutation strength adaptation realizes on the tracking problem with linear dynamics of the target. Inserting Eq. (11) shows that for that mutation strength,

$$\sigma \simeq \frac{\sqrt{2}\delta}{c_{\mu/\mu,\lambda}}, \tag{20}$$

holds, and that the resulting distance between the population centroid and the target is

$$R \simeq \frac{\sqrt{2}N\delta}{\mu c^2_{\mu/\mu,\lambda}}. \tag{21}$$

Figure 5 illustrates the quality of the approximation. While again for $N = 40$ some deviations can be observed, the agreement is excellent for $N = 400$. Comparison with the results Eqs. (12) and (13) for the optimal mutation strength derived in Section 4 shows that the mutation strength that cumulative step length adaptation realizes is by a factor of $\sqrt{2/3} \approx 0.82$ smaller than what it would ideally be, and that the distance between the population centroid and the target is by a factor of $4\sqrt{2}/(3\sqrt{3}) \approx 1.09$ larger than it would be if the optimal mutation strength were achieved.

6 Discussion and Conclusions

In this paper, the performance of the $(\mu/\mu, \lambda)$-ES with cumulative step length adaptation has been analyzed on the tracking problem with linear dynamics of the target. It has been seen that the adaptation of the step length works in the sense that a mutation strength that ensures that the target can be tracked is realized. The mutation strength in Eq. (20) exceeds the minimum mutation strength $\sigma = \delta/c_{\mu/\mu,\lambda}$ required to keep pace with the target that was obtained in Section 4. However, it has also been seen that the mutation strength that is realized is below the mutation strength given in Eq. (12) that was found to be optimal. An explanation for that behavior can be found from Eq. (20). The mutation strength that cumulative step length adaptation realizes on the tracking problem with linear dynamics of the target is the same as that realized on the static sphere. While the linearly dynamic target ideally calls for a larger step size, the strategy cannot anticipate changes of the objective and realizes the mutation strength that is statically optimal. The step length that is realized is thus probably the best that one can hope for. Our concern voiced in (Arnold and Beyer, 2002b) that the correlations in the sequence of steps taken that result from the linear motion of the target could lead to a degradation of the performance of cumulative step length adaptation has proven groundless at least in the limit of infinite search space dimensionality.

It has furthermore been seen that the distance at which the target is tracked depends on the search space dimensionality, the speed of the target, and the population size parameters that the strategy employs. In order to minimize that distance, the population size parameters must be chosen such that the denominator of the fraction in Eq. (21) is maximized. Interestingly, that denominator equals the progress rate of the $(\mu/\mu,\lambda)$-ES on the static sphere derived in (Beyer, 2001; Rechenberg, 1994). Recommendations with respect to the choice of population sizes that have been made for the static sphere thus also hold for the dynamic tracking problem. As already noted in (Arnold and Beyer, 2002b), the distance to the target is reduced substantially by the presence of the factor μ in the denominator of the fraction in Eq. (21) that stems from the presence of genetic repair. We therefore cannot confirm the observation by (Salomon and Eggenberger, 1997) that recombination is not beneficial for tracking a moving target.

To conclude, our findings suggest that evolution strategies with cumulative step length adaptation are well capable of handling dynamic optimization problems that amount to the task of tracking a moving target. While the use of populations and recombination does not improve the behavior of the strategy qualitatively, it does have the effect of the target being tracked more closely as the $(1,\lambda)$-ES would. Cumulative step length adaptation has turned out to generate optimal step lengths in the case of a randomly moving target and mutation strengths that are near optimal if the motion of the target is linear. Altogether, cumulative step length adaptation seems to hold promise far beyond the static situations that it was designed for.

Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft (DFG) under grant Be1578/6-3 and as part of the Collaborative Research Center *Computational Intelligence* (SFB 531) located at the University of Dortmund.

References

Angeline, P. J. (1997). Tracking extrema in dynamic environments. In *Proc. of the Sixth International Conference on Evolutionary Programming*, pages 335–345. Springer Verlag, Berlin.

Arnold, D. V. (2002). *Noisy Optimization with Evolution Strategies*. Genetic Algorithms and Evolutionary Computation Series. Kluwer Academic Publishers, Norwell, MA.

Arnold, D. V. and Beyer, H.-G. (2002a). Evolutionary optimization with cumulative step length adaptation — a performance analysis. Technical Report CI137/02, University of Dortmund, SFB 531. Submitted for publication.

Arnold, D. V. and Beyer, H.-G. (2002b). Random dynamics optimum tracking with evolution strategies. In et al., J. J. M., editor, *Parallel Problem Solving from Nature — PPSN VII*, pages 3–12. Springer Verlag, Berlin.

Bäck, T. (1996). *Evolutionary Algorithms in Theory and Practice*. Oxford University Press, New York.

Bäck, T. (1998). On the behavior of evolutionary algorithms in dynamic environments. In *Proc. of the 1998 International Conference on Evolutionary Computation*, pages 446–451. IEEE Press, Piscataway, NJ.

Beyer, H.-G. (2001). *The Theory of Evolution Strategies*. Natural Computing Series. Springer Verlag, Berlin.

Beyer, H.-G., Schwefel, H.-P., and Wegener, I. (2002). How to analyze evolutionary algorithms. *Theoretical Computer Science*, 287(1):101–130.

Branke, J. (2001). *Evolutionary Optimization in Dynamic Environments*. Genetic Algorithms and Evolutionary Computation Series. Kluwer Academic Publishers, Norwell, MA.

Bürger, R. (2000). *The Mathematical Theory of Selection, Recombination, and Mutation*. John Wiley & Sons, Chichester.

Droste, S. (2002). Analysis of the $(1+1)$ EA for a dynamically changing objective function. In *Proc. of the 2002 Congress on Evolutionary Computation*. IEEE Press, Piscataway, NJ.

Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley, Reading, MA.

Hansen, N. (1998). *Verallgemeinerte individuelle Schrittweitenregelung in der Evolutionsstrategie*. Mensch & Buch Verlag, Berlin.

Hansen, N. and Ostermeier, A. (2001). Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2):159–195.

Herdy, M. (1992). Reproductive isolation as a strategy parameter in hierarchically organized evolution strategies. In Männer, R. and Manderick, B., editors, *Parallel Problem Solving from Nature — PPSN II*, pages 207–217. Elsevier, Amsterdam.

Michalski, R. S. (2000). Learnable evolution model: Evolutionary processes guided by machine learning. *Machine Learning*, 38(1-2):9–40.

Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA.

Ravisé, C. and Sebag, M. (1996). An advanced evolution should not repeat its past errors. In Saitta, L., editor, *Proc. of the 13th International Conference on Machine Learning*, pages 400–408. Morgan Kaufmann, San Francisco, CA.

Rechenberg, I. (1994). *Evolutionsstrategie '94*. Frommann-Holzboog, Stuttgart.

Salomon, R. and Eggenberger, P. (1997). Adaptation on the evolutionary time scale: A working hypothesis and basic experiments. In *Proc. of the Third Conference on Artificial Evolution*, pages 251–262. Springer Verlag, Berlin.

Schwefel, H.-P. (1995). *Evolution and Optimum Seeking*. Wiley, New York.

Sebag, M., Schoenauer, M., and Ravisé, C. (1997). Inductive learning of mutation step-size in evolutionary parameter optimization. In *Proc. of the Sixth Annual Conference on Evolutionary Programming*, pages 247–261. Springer Verlag, Berlin.

van Nimwegen, E. and Crutchfield, J. P. (2001). Optimizing epochal evolutionary search: Population-size dependent theory. *Machine Learning*, 45(1):77–114.

Weicker, K. and Weicker, N. (1999). On evolution strategy optimization in dynamic environments. In *Proc. of the 1999 Congress on Evolutionary Computation*, pages 2039–2046. IEEE Press, Piscataway, NJ.