Performance Analysis of Derandomized Evolution Strategies in Quantum Control Experiments

Ofer M. Shir
Natural Computing Group, Leiden University
Niels Bohrweg 1, Leiden, The Netherlands
oshir@liacs.nl

Jonathan Roslund
Department of Chemistry, Princeton University
Princeton NJ, 08544, USA
jroslund@Princeton.EDU

Thomas Bäck*
Natural Computing Group, Leiden University
Niels Bohrweg 1, Leiden, The Netherlands
baeck@liacs.nl

Herschel Rabitz
Department of Chemistry, Princeton University
Princeton NJ, 08544, USA
hrabitz@Princeton.EDU

ABSTRACT

Genetic Algorithms (GAs) are historically the most commonly used optimization method in Quantum Control (QC) experiments. We transfer specific Derandomized Evolution Strategies (DES) that have performed well on noise-free theoretical Quantum Control calculations, including the Covariance Matrix Adaptation (CMA-ES) algorithm, into the noisy environment of Quantum Control experiments. We study the performance of these DES variants in laboratory experiments, and reveal the underlying strategy dynamics of first- versus second-order landscape information.

It is experimentally observed that global maxima of the given QC landscapes are located when only first-order information is used during the search. We report on the disruptive effects to which DES are exposed in these experiments, and study covariance matrix learning in noisy versus noise-free environments. Finally, we examine the characteristic behavior of the algorithms on the given landscapes, and draw some conclusions regarding the use of DES in QC laboratory experiments.

Categories and Subject Descriptors

I.2.8 [**Computing Methodologies**]: ARTIFICIAL INTELLIGENCE—*Problem Solving, Control Methods, and Search*

General Terms

Algorithms, Experimentation, Performance

*NuTech Solutions
Martin-Schmeisser-Weg 15
44227 Dortmund, Germany.

GECCO'08, July 12–16, 2008, Atlanta, Georgia, USA.

Keywords

Derandomized Evolution Strategies, Experimental Quantum Control, Laser Pulse Shaping, CMA-ES

1. INTRODUCTION

The advent of modern Evolution Strategies (ES) [3], also known as Derandomized Evolution Strategies [8], allows successful global optimization with minimal requirements concerning exogenous parameters, mostly without recombination, and with a low number of function evaluations. Furthermore, investigation of the actual strategy parameter space, such as the evolving covariance matrix, may allow for inferring properties of the underlying search landscape.

Quantum Control (QC) [22, 15, 23, 11], sometimes referred to as Optimal Control or Coherent Control, aims at altering the course of quantum dynamics phenomena for specific target realizations. There are two main threads within Quantum Control, *theoretical* and *experimental* control.

Genetic Algorithms (GAs) are the most common optimization routines in QC laboratories, likely due to historical reasons [10]. ES are speculated to perform well on QC optimization problems that possess continuous high dimensional landscapes. This is due to their specific variation operators, the self-adaptation of their mutation distribution as well as to their high performance in continuous global optimization in comparison to other methods on benchmark problems [2].

The goal of this work is to transfer specific derandomized ES variants, that had been applied in the past to noise-free QC landscapes, into QC experiments in the laboratory. The study then presents experimental observation of ES performance, and focuses on a comparison between first-order and second-order derandomized ES, that correspond to employing a number of strategy parameters scaling *linearly* (individual step-sizes) versus *quadratically* (arbitrary normal mutations, by means of a full covariance matrix) with the search space dimensionality. As a reference, we examine the performance behavior of the traditional GA, as typically employed in the laboratory. We discuss specific pulse shaping aspects that require treatment upon the application of unrestricted ES. We then analyze the performance behavior, and conduct additional simulations in order to investigate some special features of the CMA, as observed in the laboratory experiments.

The paper is organized as follows. Section 2 lays out the framework of Quantum Control and presents the specific optimization problems studied. Section 3 reviews the algorithms used in this study. Section 4 presents the experimental procedure and results, followed by an analysis and discussion in section 5. Finally, section 6 summarizes this work and discusses possible future directions.

2. FRAMEWORK: QUANTUM CONTROL

Quantum Control Theory (QCT) [14, 18] aims at manipulating the quantum dynamics of a *simulated system* by means of an external temporal control field, often arising from a laser source. The objective to be met in this control process is defined by means of a given physical observable, whose yield is subject to maximization. A quantum control landscape is thus defined as the functional dependence of the physical observable on the control variables, and may be visualized as a surface over the space of all possible controls. A simulated QCT system is typically numerically modeled with no physical constraints on the control field. Hence, such systems play an important role in theoretical calculations, but do not necessarily correspond to feasible real-world systems or controls.

Quantum Control Experiments (QCE) [11, 23], on the other hand, implement the ideas of Quantum Control in the laboratory and strive to alter the course of quantum dynamics for specific target-applications by utilizing a closed-loop learning algorithm. Here, the control yield, or success-rate, is obtained by a physical measurement of the target application. The nature of the optimization is fundamentally different than in QCT, due to the posed experimental constraints.

Rabitz [10] introduced the important concept of *feedback control*, where phase-, amplitude- and/or polarization shaping of a control field are subject to a *closed learning loop* in order to guide the quantum system toward a desired final state. This approach has been successfully applied in numerous applications and has become the common practical experimental routine in the field.

From the optimization perspective, GAs have become the most popular algorithmic component in the QCE learning loop (see, e.g., [11]), likely due to historical reasons. We would like to mention, however, two studies [24, 5] that applied Evolution Strategies to QCE, and explored a specific QC system both in experiments and simulations. The latter studies concluded that the employed Evolution Strategies were promising optimization routines.

2.1 Quantum Control Landscapes

In the past twenty years, QCT has revealed remarkable properties about its search landscapes. Among these results is that for a controllable quantum system, there is always a trap-free pathway to the top of the control landscape from any initial point, which allows locating the global maximum with first-order (gradient) information (see, e.g., [9]). This result is valid for an unconstrained control field.

The realization of a quantum system in the laboratory, however, inevitably places constraints on the quantum dynamics, and may lead to a tortuous path over the nominally trap-free landscape. Such experimental constraints are limited bandwidth and fluence, control resolution, proper basis, etc. One of the goals of this paper is to answer the question whether the equivalent to first-order information in stochastic algorithms is also sufficient for optimizing QCE landscapes in the laboratory.

2.2 Two-Photon Processes

The field of *nonlinear optics* describes optical phenomena which are observed when high intensity light passes through media. The nonlinearity is due to the interaction between the light, typically a laser field, and a dielectric media, whose field-induced polarization responds non-linearly to the incident electric field. The field of nonlinear optics offers a variety of popular Quantum Control applications. Second-order variants, which correspond to two-photon processes, are particularly attractive because of their easy implementation in the laboratory, as well as their known mathematical formulation.

2.2.1 Control Definition

The control function in spectral modulation consists of the spectral amplitude $A(\omega)$ and phase $\phi(\omega)$ functions. Most Quantum Control processes are sensitive to the phase, and phase-only shaping is typically sufficient for attaining optimal control. Our experiments include phase modulation only, where the spectral function $A(\omega)$ is fixed. The latter is well-approximated by a Gaussian and determines the bandwidth, or the pulse duration, accordingly. Note that shaping the field with phase-only modulation guarantees conservation of the pulse energy.

The spectral phase $\phi(\omega)$ is defined at n frequencies $\{\omega_i\}_{i=1}^n$ that are equally distributed across the bandwidth of the pulse. These n values $\{\phi(\omega_i)\}_{i=1}^n$ correspond to the n pixels of the pulse shaper and are the decision parameters optimized in the experimental learning loop:

$$\phi(\omega) = (\phi(\omega_1), \phi(\omega_2), ..., \phi(\omega_n)) \,. \tag{1}$$

Given a control phase, $\phi(\omega)$, one may write the spectral field of the control $E_1(\omega)$ and second-harmonic $E_2(\omega)$ pulses:

$$\begin{aligned} &E_1(\omega) = A(\omega)\exp[i\phi(\omega)] \\ &E_2(\omega) = E_1(\omega) * E_1(\omega) = \int_{-\infty}^{\infty} E_1(\omega') \cdot E_1(\omega - \omega')d\omega' \\ &I_2(\omega) = |E_2(\omega)|^2, \end{aligned} \tag{2}$$

where $I_2(\omega)$ is the spectral intensity of the second harmonic field.

2.2.2 Second Harmonic Generation (Total-SHG)

Second harmonic generation (SHG) or *frequency doubling* is a two-photon process in which an electric field interacts nonlinearly with a material and generates a single output photon with double the energy of two input photons. The total energy of this up-converted light is proportional to the integrated spectral intensity $I_2(\omega)$ of the second harmonic field.

The yield, subject to maximization, is defined as follows:

$$S_t = \int_{-\infty}^{\infty} I_2(\omega)d\omega. \tag{3}$$

The SHG signal S_t is maximized by any linear phase function of frequency, and in particular by a constant phase, i.e., spectral phases of the form $\phi(\omega) = \alpha + \beta\omega$. The SHG optimization problem has been widely investigated in [16], and it is known to have a landscape rich in structure [16].

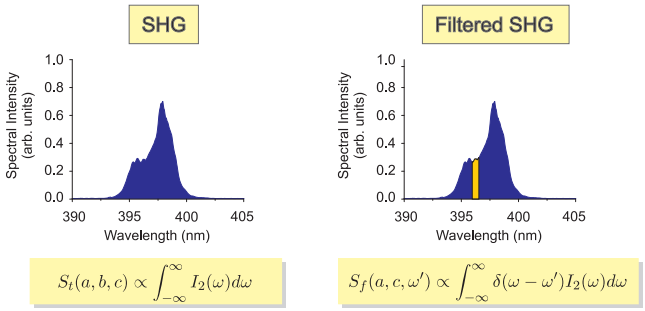

Figure 1: A spectral illustration for the Total-SHG (left) versus the Filtered-SHG (right) signals. While the Total-SHG integrates the spectral energy over the entire spectrum, the Filtered-SHG considers the integration in a specific narrow spectral window.

2.2.3 Filtered-SHG

We consider another second-order quantum optical system, which is a *filtered* variation of the SHG system. It corresponds to a *two photon absorption* (TPA) process, whose model describes, within the limits of second-order time dependent perturbation theory, the probability of making a transition from a system's ground state $|g\rangle$ to an excited state $|e\rangle$, upon activation of the laser field. Thus, a specific *transition frequency* is considered here, ω_{eg}, which practically filters the signal by means of the Dirac delta function:

$$S_f(\omega_{eg}) = \int_{-\infty}^{\infty} \delta(\omega_{eg}-\omega)\, I_2(\omega)d\omega = \\ = \left|\int_{-\infty}^{\infty} E(\omega)E(\omega_{eg}-\omega)\, d\omega\right|^2. \quad (4)$$

The filtered-SHG signal is maximized by any spectral phase that is antisymmetric about $\frac{\omega_{eg}}{2}$, i.e., spectral phases of the form $\phi(\frac{\omega_{eg}}{2}-\omega) = -\phi(\frac{\omega_{eg}}{2}+\omega)$.

Figure 1 provides an illustrative comparison between the two second harmonic generation variants considered here.

2.2.4 Problem Difficulty: Numerical Assessment

In order to assess the optimization difficulty of these maximization problems, we considered numerical simulations of the two SHG variants and conducted a simple statistical test. It is similar to the statistical test reported in [19], which only considered the total-SHG case and employed slightly different numerical modeling. The numerical test considered here pixelizes phase functions with $n = 64$ function values, which are randomly initialized in the interval $[0, 2\pi]^{64}$. We then gradually transformed the given random phases into a zero-phase by two different routines: (1) setting pixel values to zero with consistent indexing from right to left, and (2) setting function values to zero by random permutation of indices, with no repetition. Both routines eventually obtain zero-phases, which attain the maximal yield of 1.0 for both SHG problem variants. These routines were run 100 times, i.e., for 100 randomly initialized trial phases, while the yield values were recorded at each index-step per routine. Figure 2 presents typical runs for the two routines when applied to both SHG problem variants. It is observed in these plots that approximately 50% of the function values must be set to zero in order to enhance the yield value, for all cases. Once this threshold is exceeded, the yield value increases until it reaches a value of 1.0, while the profile is not necessarily monotonic. The actual profiles of the different "phasing-up" routines differ. More variables are required to be set to zero in the random indexing routine, in comparison to the consistent indexing. This is due to the shape of the weighting function (i.e., a Gaussian), which limits the contribution to the yield value from pixels which are not in the proximity of the central frequency.

This statistical test reveals that the SHG problems under investigation are non-separable upon following the formal definition.

3. ALGORITHMS

On the basis of earlier work that examined ES performance on Quantum Control landscapes [21], we restrict our study to specific ES state-of-the-art algorithms. We chose to include the traditional GA as a reference algorithm, due to the fact that it is commonly employed in QC laboratories.

The careful reader should note that we consider here first- or second-order variation information in a stochastic evolutionary algorithm (DES) as equivalent to first- or second-order search information of a deterministic algorithm.

3.1 First-Order DES: The DR2 Algorithm

The second derandomized ES variant [13] aims to accumulate information about the correlation or anti-correlation of past mutation vectors in order to adapt the *global step-size* as well as the *individual step-sizes* by introducing a quasi-memory vector. This accumulated information allows for omitting the stochastic element in the adaptation of the strategy parameters - updating the strategy parameters only by means of successful past mutations, rather than with random steps. This method stores first-order information by means of its $\mathcal{O}(n)$ strategy parameters, where n is the search space dimensionality.

We follow the recommended population size for $(1, \lambda)$ derandomized ES (see, e.g., [12]), and set $\lambda = 10$. By definition, this strategy does not apply recombination. The mutation step for the k^{th} individual, $k = 1 \ldots \lambda$, reads:

$$\vec{x}^{(g+1)} = \vec{x}^{(g)} + \delta^{(g)}\vec{\delta}^{(g)}_{scal}\vec{z}_k \qquad \vec{z}_k \sim \vec{\mathcal{N}}(0,1), \quad (5)$$

where the evolving strategy parameters $\delta^{(g)}$ and $\vec{\delta}^{(g)}_{scal}$ are the *global step-size* and the *variation directional vector*, respectively. We refer the reader to [13] for the detailed description of the defining update steps of the strategy parameters.

It should be noted that this specific DES variant has shown particularly excellent behavior on a specific QC noise-free problem, namely the dynamic molecular alignment (see, e.g., [20]). Given the default parameterization of the alignment problem, the DR2 algorithm outperformed other algorithms, including the CMA-ES (the algorithm to be discussed in the following section), and always obtained exclusively the optimal family of solutions. That numerical observation provides us with the motivation to employ it in our QC experiments.

3.2 Second-Order DES: The (μ_W, λ) CMA

The (μ_W, λ)**-CMA-ES** algorithm [8, 7] applies *principal component analysis* (PCA) to the *selected* mutations during the evolution, also referred to as *"the evolution path"*, for the adaptation of the covariance matrix of the distribution. The concept of *weighted recombination* is introduced: applying intermediate multi-recombination to the best μ out of λ offspring with given weights $\{w_i\}_{i=1}^{\mu}$. The result is denoted

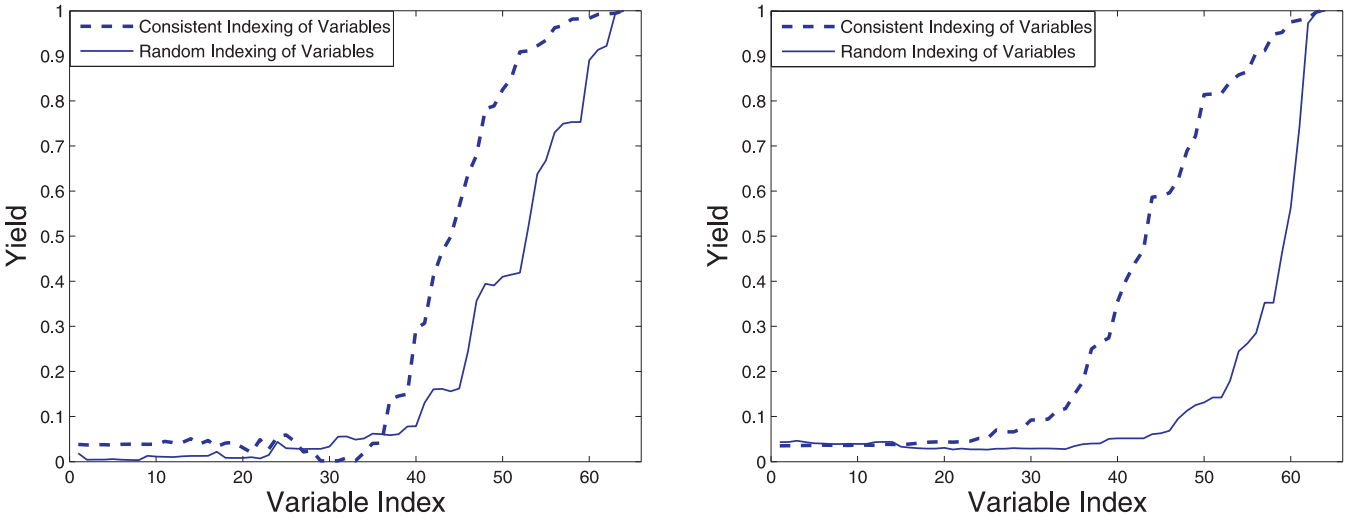

Figure 2: Transforming randomly-initialized phases into a zero-phase, pixel-by-pixel, either by (1) consistently indexing the phase function from right to left, or (2) by randomly selecting phase function indices, without repetition. The attained yield per index-step is recorded at every time step. Typical runs are presented for the two routines applied to the SHG problem variants. Left: Filtered-SHG system; Right: Total-SHG system.

with $\langle \vec{x} \rangle_W$. $\mathbf{C}^{(g)} \in \mathbb{R}^{n \times n}$ is the covariance matrix, which is eigen-decomposed as $\mathbf{C}^{(g)} = \mathbf{B}^{(g)}\mathbf{D}^{(g)}\left(\mathbf{B}^{(g)}\mathbf{D}^{(g)}\right)^T$. By definition, this method stores second-order information by means of its $\mathcal{O}(n^2)$ strategy parameters. The characteristic generation step for the k^{th} individual, $k = 1, \ldots, \lambda$, reads:

$$\vec{x}_k^{(g+1)} = \langle \vec{x} \rangle_W^{(g)} + \sigma^{(g)}\mathbf{B}^{(g)}\mathbf{D}^{(g)}\vec{z}_k^{(g+1)} \tag{6}$$

We refer the reader to [7] for the detailed description of the defining update steps of the strategy parameters. We set the population size to the recommended values for $n = 64$, which is the number of pixels used by the pulse shaper in these experiments: $\mu = 8$ and $\lambda = 16$.

3.3 Traditional GA

We use the traditional GA [6], with bitstring representation of $l = 6$ bit resolution per pixel. It employs a fixed population of $\mu = 30$ individuals. The mutation rate for a bit-flip is $p_m = 0.005$, and the selection mechanism is based on the better half, while the single best individual is always kept (*elitism*). These parameters were collectively optimized to allow sufficient resolution so as to arrive at the highest quality solution with the fastest convergence.

4. EXPERIMENTAL PROCEDURE

We describe here the experimental procedure, where we applied the specified algorithms to the laboratory quantum control systems.

All algorithms were coded in LabView[1]. Regarding the technical specification of the laboratory, the laser source is a Ti:sapphire femtosecond system with 1.8 mJ pulses generated at 1 kHz. The pulses are centered at ~ 800 nm, with a bandwidth of $\Delta\lambda \sim 10$ nm, yielding $\Delta\tau \sim 100$ fs pulses, FWHM. The employed *spatial light modulator* (SLM) consists of 128 liquid crystal pixels, but the experiments used only 64 pixels by coupling together pairs of adjacent pixels. Thus, here we set $n = 64$.

For total SHG S_t, the amplified pulses are delivered to a 100 μm type-I BBO crystal, and the time integrated SHG signal is recorded with a photodiode and boxcar integrator. For the filtered SHG signal S_f, unamplified seed pulses are focused onto a 100 μm type-I BBO crystal, and the resultant up-converted light is analyzed with a spectrometer.

For the two DES variants employed, initial step sizes are set to $\frac{1}{4}$ of the object variables initialization intervals.

4.1 Preliminary ES Failure: Stretched Phases

When applied to the experimental setup, the derandomized ES variants initially suffered from pre-mature convergence to sub-optimal solutions of yield ≈ 0.75, where the maximum value is 1.0. Upon examination of the attained optimized phases in the decision space, they were always observed to be *highly sloped linear phases*. We offer the following explanation for the suboptimal convergence.

In practice, the use of a pixel-based pulse shaper represents a desired spectral phase $\phi(\omega)$ with an appropriate staircase-approximation. The result of this spectral domain pixelization is a temporal electric field of the form

$$e(t) = \sum_n \tilde{e}(t - n\tau)\text{sinc}\left(\frac{\pi t}{\tau}\right), \tag{7}$$

with $\tilde{e}(t)$ as the desired electric field, and where $\tau = \frac{1}{\Delta\nu}$ is the inverse frequency spacing per pixel. The result of pixelization is the creation of a series of replica pulses, coincidentally positioned at the zeros of the **sinc** envelope function. Application of a significant linear phase, which corresponds to a temporal shift of the pulse (see, e.g., [4]), increases the influence of the replica pulses by shifting them away from the zeros. Thus, excessively steep linear phases

[1]Code is available from the authors upon request.

Table 1: QCE Algorithms: Performance

Routine	Filtered-SHG		Total-SHG	
	Avg. Yield	0.9 Eval	Avg. Yield	0.9 Eval
GA	0.95	4665	0.95	5557
DR2	0.93	2159	0.72	NA
CMA	0.95	841	0.98	766

breakup the pulse energy into multiple parasitic replica pulses and result in suboptimal yields [17].

Phases that differ by 2π *radians* are mathematically equivalent. This periodic nature of the phase in $[0, 2\pi]$ practically poses periodic boundary conditions on the shaping modulator. Given $0 < \varepsilon \leq 2\pi$, the so-called *phase wrapping* operator is implemented as follows:

$$\begin{aligned} \phi_i = 2\pi + \varepsilon \quad &\longrightarrow \quad \tilde{\phi}_i := \varepsilon \\ \phi_j = -\varepsilon \quad &\longrightarrow \quad \tilde{\phi}_j := 2\pi - \varepsilon \end{aligned} \tag{8}$$

or simply as $\tilde{\phi}_i := \phi_i \bmod 2\pi$. From an optimization perspective, this means that the search space is an n-dimensional hypercube spanning a length of 2π in each dimension.

The GA, which is traditionally employed in the laboratory, is less likely to locate highly-steep linear phases since the $[0, 2\pi]$ boundaries are implicitly implemented by means of the *phenotypic mapping*. It seems that a search in an unrestricted domain, as employed by the ES variants in hand, is likely to irreversibly stretch the candidate phases and suffer accordingly from convergence to highly sloped linear phases with sub-optimal yield.

The implementation of *periodic boundary conditions* into the ES algorithms, by means of coupling the *phase wrapping operator* to the mutation operator, alleviated this suboptimal convergence. It should be noted that the restriction of the search domain by means of the phase wrapping operator does not in any way restrict the quality or diversity of attainable solutions.

4.2 Performance Evaluations

Table 1 presents the results of the two reported systems, averaged over 10 experiments. We consider the final yield (averaged over the last 50 iterations), as well as the number of evaluations required to cross a yield threshold of 0.90, as the performance criteria per experiment. Figure 3 presents an averaging of the runs, with attained yield as a function of the required number of function evaluations. Note that this averaging procedure takes into account all 10 runs, whereas the convergence data shown in Table 1 considers only the relevant runs that exceeded the 0.90 yield threshold. Figure 4 presents histograms for the different algorithms with final yield versus the number of runs.

As reflected from the experimental results, the CMA performed best on the given experimental systems, both in terms of final yield as well as convergence speed. We would like to emphasize the extraordinary boost of convergence speed provided by the CMA relative to the GA, which is significant in the laboratory. Moreover, the CMA has a sharp and rapid convergence profile, in contrast to the inefficient hill-climbing capability of the GA. This profile is easy to identify as there is no ambiguity about convergence, and thus it is another attractive feature for the laboratory user.

5. ANALYSIS AND DISCUSSION

We discuss here the experimental results and the algorithmic behavior.

5.1 Diversity of Solutions

As mentioned earlier in section 2.2, the filtered SHG system possess a family of nontrivial phases that correspond to global maxima. Interestingly, each run for the filtered SHG case converged to a distinct antisymmetric phase. This collection of different solutions provided a practical perspective concerning the richness of QC landscapes and their underlying level sets.

5.2 Sensitivity to Noise

The CMA-ES and the GA performed in a satisfactory manner on the given control problems and did not seem to be significantly impaired by the existence of noise in the experimental system. The DR2, on the other hand, suffered from high-sensitivity to the initial step-size: the default setting failed to obtain satisfying results, and a fine-tuning process obtained a narrow range of initial values that obtained fine results. Generally speaking, the DR2 performance was disappointing, particularly in comparison to noise-free calculations that were reported in the past [20]. A proposed explanation for this behavior could be the lack of recombination, which has been shown to be a crucial ES component in noisy environments (see, e.g., [1]).

5.3 Covariance Learning

Recording the CMA data during the optimizations allows an analysis of the evolutionary search process. It was found that the covariance matrix remains diagonal during the search, or equivalently, the CMA does not use its second-order information (i.e., *rotations*) when climbing up the landscape. This observation seems to be consistent with the QCT landscape analysis outlined in section 2.1.

Figure 5 presents a typical CMA run for the optimization of total-SHG in the laboratory and shows the yield and step-size verses function evaluations. Figure 6 presents the square-root of the covariance matrix eigenvalues as a function of the number of experiments as well as the Euclidean distances between the best phase variables of successive iterations, i.e.,

$$d^{(g+1)} = \|\vec{\phi}^{(g+1)}_{best}(\omega) - \vec{\phi}^{(g)}_{best}(\omega)\|, \tag{9}$$

where $\vec{\phi}_{best}(\omega)$ is as in Eq. 1.

We conducted an equivalent test in a **noise-free simulator** for the total-SHG problem[2]. The convergence profile on the simulator is observed to be similar to the laboratory experiment, i.e., rapid climbing-up of the landscape without using second-order information. Figure 7 presents a typical CMA run on the simulator. However, upon approaching the *top of the landscape*, one of the covariance matrix eigenvalues dramatically grows, as shown in Figure 8. This behavior was observed to be typical in all runs. The corresponding eigenvector is always a flat phase, suggesting that the CMA discovers the invariance of a constant phase on the SHG signal. The phase Euclidean trajectories are plotted as well in Figure 8, showing some minor activity during this growth stage, corresponding to super-fine tuning of the spectral phase. The yield values, nonetheless, do not seem

[2]The simulator was implemented in LabView with Lab2.

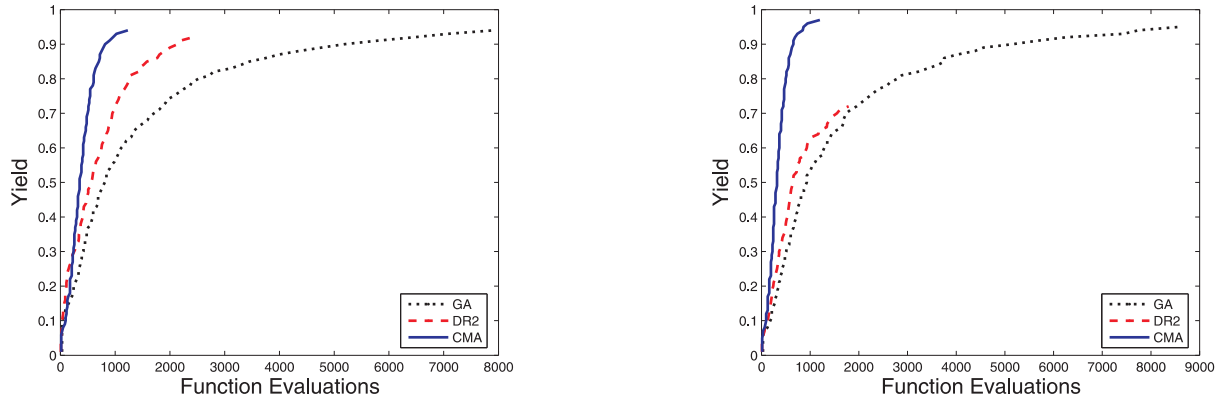

Figure 3: Averaged runs of the algorithms over 10 runs. Left: Filtered-SHG system; Right: Total-SHG system.

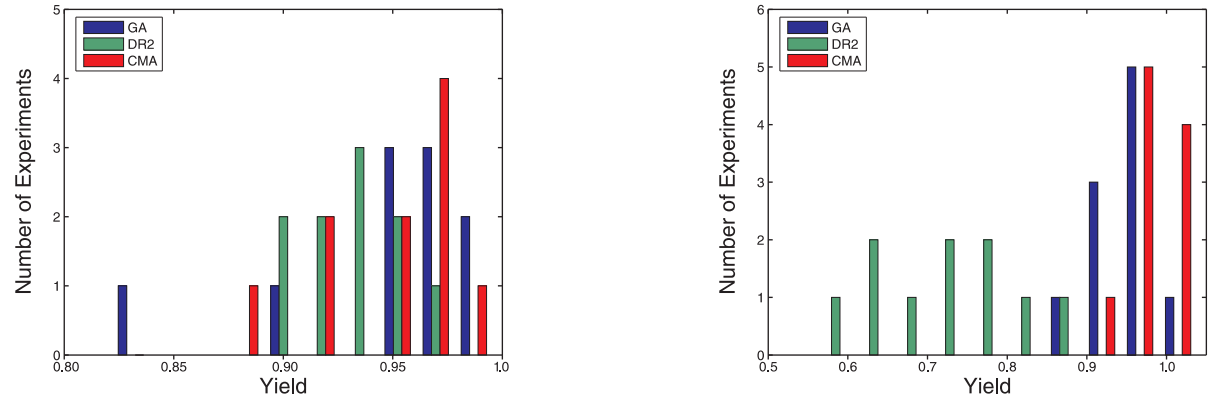

Figure 4: Success-rate (yield) histograms. Left: Filtered-SHG system; Right: Total-SHG system.

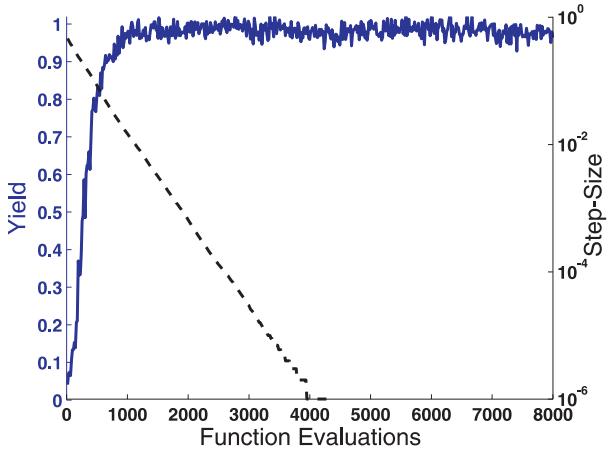

Figure 5: CMA optimization of the Total-SHG in the laboratory. Yield (solid line, left axis) and step-size (dashed line, right log-scaled axis), versus function evaluations.

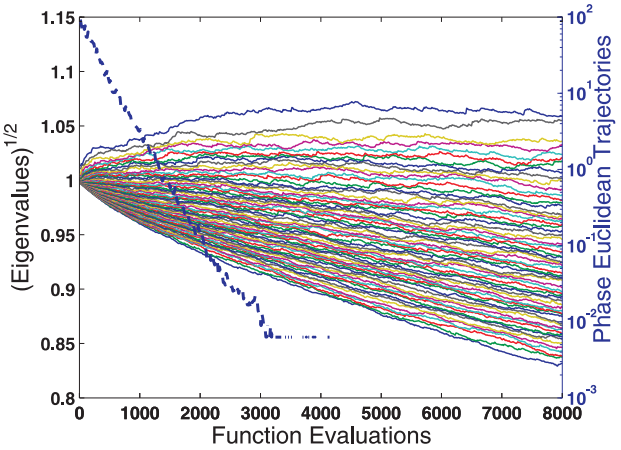

Figure 6: CMA optimization of the Total-SHG in the laboratory. Square-root of the 64 eigenvalues of the covariance matrix (solid thin lines, left axis), and phase Euclidean trajectories (bold points, right log-scaled axis), versus function evaluations. Missing trajectory points correspond to zero values.

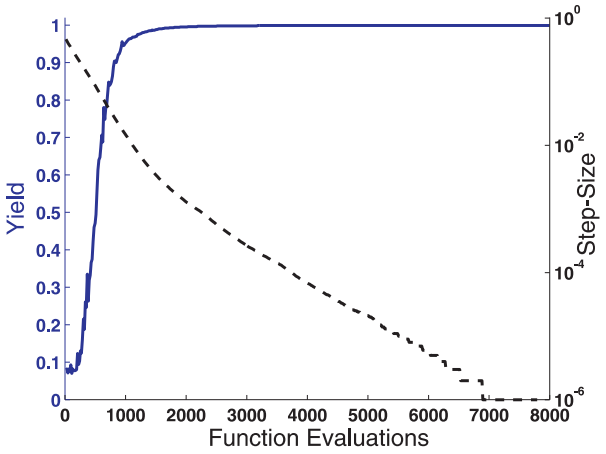

Figure 7: CMA optimization of the Total-SHG on a noise-free simulator. Yield (solid line, left axis) and step-size (dashed line, right log-scaled axis), versus function evaluations.

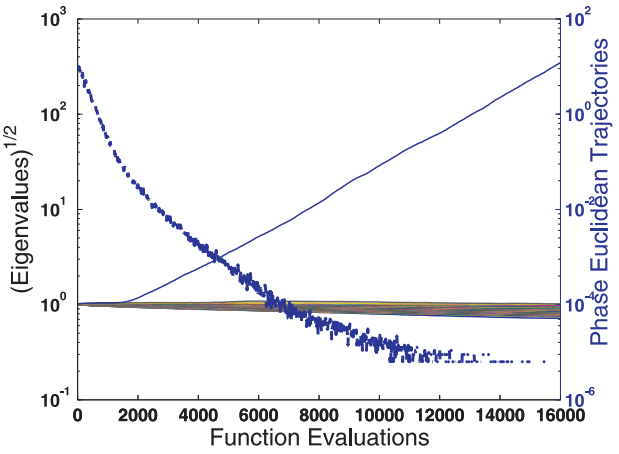

Figure 8: CMA optimization of the Total-SHG on a noise-free simulator. Square-root of the 64 eigenvalues of the covariance matrix (solid thin lines, left log-scaled axis), and phase Euclidean trajectories (bold points, right log-scaled axis), versus function evaluations. Missing trajectory points correspond to zero values. The single exploding eigenvalue can easily be identified in this scale.

to be further improved during this process, at least in the precision available. In practice, the parameter adaptation during this fine-tuning stage produces fitness variations below that of the system noise in the laboratory, which explains its absence in laboratory optimizations.

5.4 Simulations: Zeroth-Order CMA

Given the experimental observation reported in the previous section, we were interested in testing the CMA while removing its covariance learning components. In essence, we leave the CMA only with the step-size as a strategy parameter, and fix the covariance matrix as an identity matrix. This is a zeroth-order ES with normal mutations subject to hyperspheres as the equi-density probability surfaces. In order to assess the zeroth-order CMA behavior on the given QC systems, we conducted additional simulations with two variants of the algorithm:

- (μ_W, λ)-CMA with $\mathbf{C} = \mathbb{I}$.
- $(1, \lambda)$-CMA with $\mathbf{C} = \mathbb{I}$.

The simulations were conducted for both systems - total-SHG as well as filtered-SHG; we considered both a noise-free simulator as well as a simulator with noise.

The results of the simulations show that the CMA performance is not hampered at all on both systems when removing its covariance learning components: the (μ_W, λ)-CMA with $\mathbf{C} = \mathbb{I}$ performs as well as the original CMA, in terms of final attained yield and convergence speed. This observation is valid for noise-free as well as for noisy simulations, and has also been **confirmed by additional laboratory experiments**. However, when the weighted recombination operator was removed, the $(1, \lambda)$-CMA with $\mathbf{C} = \mathbb{I}$ did not converge, nor did it even climb-up from the initial yield at the bottom of the landscape. We thus conclude that it is possible to optimize the given simulated QC landscapes by a zeroth order ES, as long as the weighted-recombination operator is kept.

6. SUMMARY AND OUTLOOK

We presented a survey of derandomized Evolution Strategies and a Genetic Algorithm to a set of Quantum Control systems in the laboratory.

While the QC systems examined here possess easily understood global optima, the search is conducted over a highly complex, curvilinear control landscape. This can be assessed by examining the covariance matrix of the actual decision parameters of the recorded evolving control phases, in contrast to the obtained covariance matrix of the CMA-ES, which attempts to learn successful search variations.

We found that employing the ES variants with default settings unrestrictedly on the given QC landscapes resulted in pre-mature convergence to sub-optimal phases with highly sloped linear profiles. We analyzed this effect, and introduced the wrapping operator into the ES framework. The latter solved the observed problem.

The CMA-ES significantly outperformed the other algorithms in terms of final yield as well as in convergence speed. It introduced a significant increase in convergence speed to the typical performance of the GA in the laboratory and is a promising tool for future laboratory experiments. While analyzing its behavior, it was experimentally confirmed that its second-order information was not used when climbing-up the landscape. It is crucial to note that this behavior is not isolated to the SHG systems. Rather the same CMA behavior has additionally been observed during optimizations of atomic Rubidium, which relies upon strong-field, dynamic Stark shifting and exhibits a complex, nontrivial optimal solution. We also conducted *noise-free simulations* of the CMA-ES applied to the systems. The latter calculations revealed interesting behavior of the covariance matrix, upon

approaching the top of the landscape. A single eigenvalue consistently explodes with a corresponding eigenvector of flat phase. We suggest that this is due to the fact that the CMA successfully learned the invariance of a constant phase in these problems. Furthermore, we considered zeroth-order versions of the CMA in simulation, where the covariance learning component was removed. The latter performed extremely well, as long as the weighted-recombination operator was kept.

As future research, we would like to consider the following:

1. Investigating the DR2 sensitivity to the noise.
2. Extending the set of QC experimental systems, to particularly difficult problems with unknown solutions.
3. Extracting landscape info from evolutionary pathways.

Acknowledgments

Ofer Shir would like to thank Niko Hansen for the helpful correspondence and Darrell Whitley for the fruitful discussions. Thomas Bäck and Ofer Shir acknowledge support from the Dutch Foundation for Fundamental Research on Matter (FOM). Herschel Rabitz and Jonathan Roslund acknowledge support from the U.S. Department of Energy.

7. REFERENCES

[1] D. V. Arnold and H.-G. Beyer. Local Performance of the $(\mu/\mu_I, \lambda)$-ES in a Noisy Environment. In W. Martin and W. Spears, editors, *Foundations of Genetic Algorithms, 6*, pages 127–141, San Francisco, CA, 2001. Morgan Kaufmann.

[2] A. Auger and N. Hansen. Performance Evaluation of an Advanced Local Search Evolutionary Algorithm. In *Proceedings of the 2005 Congress on Evolutionary Computation CEC-2005*, pages 1777–1784, Piscataway, NJ, USA, 2005. IEEE Press.

[3] H.-G. Beyer and H.-P. Schwefel. Evolution Strategies a Comprehensive Introduction. *Natural Computing: an international journal*, 1(1):3–52, 2002.

[4] R. Bracewell. *The Fourier Transform and Its Applications*. McGraw-Hill Book Company, 1965.

[5] R. Fanciulli, L. Willmes, J. Savolainen, P. van der Walle, T. Bäck, and J. L. Herek. Evolution Strategies for Laser Pulse Compression. In *Proceedings of the International Conference Evolution Artificielle*, volume 4926 of *Lecture Notes in Computer Science*. Springer, 2008.

[6] D. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley, Reading, MA, 1989.

[7] N. Hansen and S. Kern. Evaluating the CMA Evolution Strategy on Multimodal Test Functions. In *Parallel Problem Solving from Nature - PPSN V*, volume 1498 of *Lecture Notes in Computer Science*, pages 282–291, Amsterdam, 1998. Springer.

[8] N. Hansen and A. Ostermeier. Completely Derandomized Self-Adaptation in Evolution Strategies. *Evolutionary Computation*, 9(2):159–195, 2001.

[9] T.-S. Ho and H. Rabitz. Why do Effective Quantum Controls Appear Easy to Find? *Journal of Photochemistry and Photobiology A: Chemistry*, 180(3), Jun 2006.

[10] R. S. Judson and H. Rabitz. Teaching Lasers to Control Molecules. *Phys. Rev. Lett.*, 68(10):1500–1503, Mar 1992.

[11] P. Nuernberger, G. Vogt, T. Brixner, and G. Gerber. Femtosecond Quantum Control of Molecular Dynamics in the Condensed Phase. *Phys Chem Chem Phys.*, 9(20):2470–2497, 2007.

[12] A. Ostermeier, A. Gawelczyk, and N. Hansen. A Derandomized Approach to Self Adaptation of Evolution Strategies. *Evolutionary Computation*, 2(4):369–380, 1994.

[13] A. Ostermeier, A. Gawelczyk, and N. Hansen. Step-Size Adaptation Based on Non-Local Use of Selection Information. In *Parallel Problem Solving from Nature - PPSN III*, volume 866 of *Lecture Notes in Computer Science*, pages 189–198. Springer, 1994.

[14] A. P. Peirce, M. A. Dahleh, and H. Rabitz. Optimal Control of Quantum-Mechanical Systems: Existence, Numerical Approximation, and Applications. *Phys. Rev. A*, 37(12), Jun 1988.

[15] H. Rabitz, R. de Vivie-Riedle, M. Mutzkus, and K. Kompa. Whither the Future of Controlling Quantum Phenomena? *Science*, 288:824–828, May 2000.

[16] J. Roslund, M. Roth, and H. Rabitz. Laboratory Observation of Quantum Control Level Sets. *Phys. Rev. A*, 74, 2006.

[17] M. Roth. *Optimal Dynamic Discrimination in the Laboratory*. PhD thesis, Princeton University, 2007.

[18] S. Shi and H. Rabitz. Quantum Mechanical Optimal Control of Physical Observables in Microsystems. *Chem. Phys.*, 92(364), Jan 1990.

[19] O. M. Shir and T. Bäck. The Second Harmonic Generation Case Study as a Gateway for ES to Quantum Control Problems. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-2007*, pages 713–721, New York, NY, USA, 2007. ACM Press.

[20] O. M. Shir, J. N. Kok, M. J. Vrakking, and T. Bäck. Gaining Insight into Laser Pulse Shaping by Evolution Strategies. In *IWINAC*, volume 4527 of *Lecture Notes in Computer Science*, pages 467–477. Springer, 2007.

[21] O. M. Shir, C. Siedschlag, T. Bäck, and M. J. Vrakking. Evolutionary Algorithms in the Optimization of Dynamic Molecular Alignment. In *2006 IEEE World Congress on Computational Intelligence*, pages 9817–9824. IEEE Computational Intelligence Society, 2006.

[22] W. S. Warren, H. Rabitz, and M. Dahleh. Coherent Control of Quantum Dynamics: The Dream Is Alive. *Science*, 259:1581–1589, Mar 1993.

[23] T. C. Weinacht and P. H. Bucksbaum. Using Feedback for Coherent Control of Quantum Systems. *Journal of Optics B*, 4(3), 2002.

[24] D. Zeidler, S. Frey, K.-L. Kompa, and M. Motzkus. Evolutionary Algorithms and their Application to Optimal Control Studies. *Phys. Rev. A*, 64(2):023420, Jul 2001.