Stochastic Search on Graphs and Genetic Algorithms

Matthias Löwe
Universität Bielefeld
Fakultät für Mathematik
Postfach 100131
33501 Bielefeld

Abstract

We analyze the convergence of a wide class of Genetic Algorithms by reducing them to a class of stochastic search algorithms on graphs. We show that, if the underlying random walk is irreducible and reversible, the invariant measure has a Gibbsian structure and that a logarithmic cooling schedule similar to that used in Simulated Annealing guarantees convergence to the uniform distribution on the set of minima of an energy function. Moreover, using spectral gap estimates, we bound the entrance time into typical subsets of the state space stochastically in terms of the Gibbs measure at time t.

Keywords: Genetic Algorithms, Spectral gap, Markov processes, Simulated Annealing

1 Introduction

In the sequel we are going to give a small contribution to the following well studied problem:
Let (X, E) be a finite, but very large graph with vertex set X and edge set E. Let $U : X \to \mathbb{R}$ be a real valued function of the vertices. Then the task to find the minimum value of U in reasonable time or, equivalently, to find the vertex, where this minimum is attained, is not an easy one.

The simplest deterministic search method, which definitely finds the global minimum, is the exhaustive or global search, that calculates explicitly the values of any vertex $x \in X$. For large sets X (which e.g. in the case of image processing are in the order of 10^{10^4}) this algorithm even if implemented on very fast computers obviously takes too long to be interesting. On the other hand the so called local search, an algorithm that makes use of the topological structure on X by starting in a vertex $x_0 \in X$ and then stepping to x_1, that vertex in $\{x_0\} \cup \mathcal{N}(x_0)$ with smallest U-value (where $\mathcal{N}(x)$ denotes the neighborhood of a vertex $x \in X$) and so on, soon gets trapped in a local minimum which usually is not the global one. So this algorithm is applicable only if one is satisfied with suboptimal solutions.

The class of stochastic search algorithms with the well known Simulated Annealing (SA) algorithm (introduced in 1983 by Kirkpatrick et. al. in [K83]) as its most prominent representative are based on the following idea. Run a time inhomogeneous Markov-chain or Markov-process on X such that the invariant measure at time t becomes more and more concentrated on the set of minima of U (e.g. in many cases the invariant measure is a Gibbsian measure with energy U and at some inverse temperature β_t with respect to some a priori measure on X). So if one changes the invariant measure carefully enough (in t) such that the process is able to follow its invariant measure, the distribution of the Markov chain will converge to the uniform distribution on the set of global minima of U.

The problems implicit in the description of this class of stochastic search algorithms, i.e. the convergence of the algorithm and the optimal choice of the way the invariant measure changes in time (in the case of a Gibbsian measure as above this change is governed by the way β_t converges to infinity; in this case we will refer to the functional dependence of β_t from t as cooling schedule or freezing schedule) has intensively been studied for SA processes.