

Evolutionary Computation with Simulated Annealing: Conditions for Optimal Equilibrium Distribution

Enrique C. Segura

Department of Computer Science, University of Buenos Aires
Ciudad Universitaria, Pab.I PB., Intendente Guiraldes s/n,
(1428) BuenosAires, Argentina

ABSTRACT

In this paper a thermodynamic approach is presented to the problem of convergence of evolutionary algorithms. The case of the Simulated Annealing algorithm for optimisation is considered as a simple evolution strategy with a control parameter allowing balance between the probability of obtaining an optimal or near-optimal solution and the time that the algorithm will take to reach equilibrium. This capacity is analysed and a theoretical frame is presented, stating a general condition to be fulfilled by an evolutionary algorithm in order to ensure its convergence to a global maximum of the fitness function.

Keywords: evolutionary computation, simulated annealing, thermodynamics of equilibrium, detailed balance, ergodicity.

1. INTRODUCTION

In the last years, Evolutionary Algorithms (EA's) have come to be considered as a very powerful and versatile class of techniques for a wide variety of complex optimisation tasks. Different evolution-inspired paradigms have been proposed (genetic algorithms, evolution programs, evolution strategies, genetic programming [6],[8]), each of them best suited for a particular type of problem. The efficiency of these approaches is out of question (at least in this paper). What we want to examine here is a rather theoretical side of the general evolutionary approach. This aspect is the question about the conditions of convergence of an evolution-inspired program, i.e. under which conditions such type of program can be guaranteed to attain, in a finite time, the optimal solution for a given problem. We said *rather*, but in no way *purely* theoretical, since the conclusions of this analysis could be, in certain cases, decisive for the usefulness of a technique.

Then, the first question that naturally arises is: can EA's be guaranteed to converge to a globally optimal population? And the answer is: no, unless infinite population is assumed [5],[7]. This is a quite general result, and it could even be considered irrelevant, since for most of real-life optimisation problems, a good sub-optimal solution can be

obtained with a well parameterised EA, without a too high computational cost, and this is often all what is needed.

But a second, more general question, could be of practical interest: can EA's be ensured to converge to a globally optimal population with a certain predetermined probability? (i.e. to know a priori the asymptotic probability distribution?). This is not a theoretical question anymore, since the possibility to know a priori the probability of global convergence and, moreover, the ability to control it via a proper parameter setting is a quite attractive perspective, especially for those applications in which an optimal or near-optimal solution is required.

The task of providing an answer (or answers) to the question stated above, could be considered as the subject of a very ambitious program, which would mainly include, among its theoretical concerns, the study of the conditions to be fulfilled in order for this program to be sensible and, as a practical, technical task, the design and development of efficient EA's satisfying those conditions.

Being this a previous exploratory work, we do not claim to have many specific answers. Nevertheless, we can bring a concrete example of a stochastic technique that, despite its different original inspiration, can legitimately be considered as a simple evolution strategy and, hence, a member of the EA family in its own right. We are talking about the Simulated Annealing (SA) algorithm.

2. THE SA ALGORITHM

The SA algorithm is a computational formalism inspired in the thermodynamics of equilibrium for a physical system with a large number of particles, but fixed in its size, so that it can only exchange energy with the environment. This model is known in statistical thermodynamics as the *canonical ensemble*. In more practical terms, the SA tries to emulate the physical process by which a crystal configuration can be obtained by melting some piece of material and then cooling it down very slowly. This is a common technique for manufacturing fine crystal goods. If cooling is too fast, a *glass*, i.e. imperfect crystal is obtained, associated with

metastable configurations, local minima for the energy of the physical system. On the other hand, a perfect crystalline structure, associated with a global minimum of the energy, requires slow cooling (until freezing) to be obtained.

Obviously, the SA is suitable either for minimisation or maximisation problems. Fig. 1 summarises the logic flow of a program implementing the SA algorithm for the problem of maximising a function F taking real values over a set S . C is the *candidate generation distribution*, i.e. $C(r,s)$ represents the probability of generating r as a candidate when the current solution produced by the algorithm is s . Similarly, A , the *acceptance probability distribution*, that is, the probability, at time t , of accepting candidate r when the system is in s , can be expressed in a compact manner as $\min\{1, \exp(\mathbf{DF}/T(t))\}$, being $\mathbf{DF} = F(r) - F(s)$. Finally, a stochastic matrix P , representing the transition probability distribution of the whole process, will result from the combination of both C and A . In other words, the probability for the algorithm to change from state s to state r is the probability to generate r as a candidate when the system is in s , times the probability of accepting it.

The more relevant results concerning convergence of the SA algorithm can be summarised as follows [9],[11]: assuming that

i) C is symmetric:

$$C(r,s) = C(s,r) \quad \forall r,s \in S \quad (1)$$

SIMULATED ANNEALING

```

-  $t=0$ 

- Choose initial  $s(0)$  from  $S$  domain of configurations

- While (~ Termination Condition)
  {
     $t++$ 

    Generate candidate  $c$  from  $C(r,s)$ 

    If  $F(c) > F(s(t-1))$ 
       $s(t) = c;$ 
    Else
      {
        generate  $\mathbf{a} \gg U[0,1]$ 
        if  $\exp(\mathbf{DF}/T(t)) > \mathbf{a}$ 
           $s(t) = c;$ 
        else
           $s(t) = s(t-1);$ 
      }
  }
    
```

Fig. 1. Basic pseudo-code for the Simulated Annealing algorithm.

ii) C satisfies the *Markov property*:

$$C(s(t+1)|s(t)) = C(s(t+1)|s(t),s(t-1),...) \quad (2)$$

i.e. the probability of generating a candidate is independent of the past states of the system.

iii) P is *ergodic*, that is, without any transient subsets. In other words, any state may be re-visited after a long enough time (which implies that P has a left eigenvector \mathbf{p} with eigenvalue 1).

Then the following holds:

a) Equilibrium (for fixed T)

$$\lim_{t \rightarrow \infty} P(s(t) = w | s(0) = \mathbf{h}) = \mathbf{p}(w) \propto \exp\left(\frac{F(w)}{kT}\right) \quad (3)$$

where π is the *Gibbs* or *Boltzmann* distribution and the constant k is assumed to be 1.

b) Annealing (decreasing T , equilibrium when $T=0$).

If

$$T(t) > \frac{D}{1 + \log(1+t)}$$

then

$$\lim_{t \rightarrow \infty} P(s(t) = w | s(0) = \mathbf{h}) = U[\Omega] \quad (4)$$

being D some positive constant and $U[\mathbf{W}]$ the uniform probability distribution over the set of global maxima for F . $U[\mathbf{W}]$ coincides with the Gibbs distribution for $T = 0$.

Briefly speaking, a) states that, the lower T , the higher the probability of the system to be in a global maximum after equilibrium has been reached (no matter the initial state of the system), while b) means that, provided T decreases slowly enough, the final equilibrium distribution only assigns probabilities greater than zero to global optima.

In other words, a parameter such as T in the case of SA would be a very useful feature to be provided to any EA. This is not impossible in principle, and here we have an example, since SA can be easily showed to be but an evolution strategy (ES) [10],[12] combined with a non-stationary Markov process. Replacing the term *offspring operator* by *candidate generation distribution* and *selection operator* by *acceptance probability distribution*, the SA depicted in Fig. 1 looks much like an ES with a two-member population (a $(I+I)$ -ES in the usual notation [2],[8]) and an increasing selective pressure, represented by the parameter T . If, in addition, we fix T at a constant value, the analogy is

even greater (Fig. 2). The difference is, of course, that in a classical ES (basically applied to continuous optimisation problems) the candidate generation distribution is Gaussian and the selective pressure is uniformly maximum along the whole process (i.e. the offspring replaces its parent if and only if it has a greater fitness). It is worth mentioning that for this type of EA's there exist theoretical results concerning convergence to the optimum [3].

What is, then, the essential feature that makes the SA, even at a fixed temperature, a *special case* of ES? It is the *ergodicity* associated to the candidate generation distribution (its "offspring operator" in the terminology of EA's). In other words, the Markov chain associated to that distribution has no *transient subsets*, which implies that any state of the space of possible configurations ("genotypes") has non-zero probability of being generated as a candidate, no matter the time elapsed.

3. THERMODYNAMICS OF EQUILIBRIUM AND EVOLUTIONARY ALGORITHMS

It is a well known fact from thermodynamics that the properties of a system in equilibrium are completely determined by its energy. This is the case for a canonical ensemble: if we know the energy associated to each state of the system, then its properties in equilibrium at temperature T can be computed as if we had an ensemble of identical systems, where the probability of finding one of them in any of the possible states is proportional to $\exp(-E/kT)$, the Gibbs or canonical distribution (k stands for the Boltzmann constant). In other words, given the probability distribution for the states, the mean value for any observable quantity is predetermined [1].

Therefore, allowing the system to evolve for a long enough time, at a fixed temperature, the time average of any quantity will equal the average over many systems appearing in different states with a relative frequency proportional to $\exp(-E/kT)$.

On the other hand, in an EA the basic structure of the elements in the state space is given by the *genotype*. Each of them has associated a fitness function F , that the algorithm is expected to maximise upon the evolution process. Then the system can be viewed, to some extent, as an ensemble of identically distributed genotypes (population), except, probably, for the crossover operator. If the parameters of the algorithm are fixed, the system will eventually reach its equilibrium distribution. The success of an EA resides, then, in how large is the probability assigned by this equilibrium distribution to the maxima of the fitness function F . In other words, how likely is to

find one or more of the individuals at or near a global optimum of F .

Now, for the case of SA, we know from section 2, Eq. (3) that, at fixed temperature, its equilibrium distribution is the Gibbs canonical distribution. Hence, SA is an ES with an equilibrium distribution that assigns increasing probability to global optima as the parameter T is lowered. Therefore, at equilibrium for a given T , it is easy to compute the probability to get a final state with value of F greater than an arbitrary L :

$$Z^{-1} \int_{s \in S / F(s) > L} \exp(F(s)/T) ds \tag{5}$$

where Z is the so called *partition function* (a normalisation constant for our purposes). For the case of a discrete configuration space the integration sign is to be replaced by a sum.

SIMULATED ANNEALING (finite temperature)

```

- t=0

- Choose initial s(0) from S set of possible solutions

- While (t < maximum number of generations)
  {
    t++

    /* Offspring */
    Generate offspring c from operator C(r,s)

    /* Selection */
    If F(c) > F(s(t-1))
      s(t) = c;
    Else
      {
        generate a » U[0,1]
        if exp(DF/T) > a
          s(t) = c;
        else
          s(t) = s(t-1);
      }
  }
    
```

Fig. 2. The SA as a two-member evolution strategy.

This useful feature of SA is due, of course, to the transition probability P , which is in its turn determined by both the candidate generation distribution C and the acceptance probability A (see section 2) or, in evolutionary computing terms, by the offspring and the selection operators. Two interesting questions arise: i) would it be possible to find other EA's whose genetic operators can be guaranteed to bring, with a high probability, the system (i.e. the population) into an optimal or near-

optimal fitness value? and ii) could these operators be parameterised so as to control that probability?

Generally speaking, the lower the temperature at which the system is fixed, the longer it will take to reach equilibrium. And this is the reason why question ii) is worth to be asked. There is, unavoidably, a trade-off between efficiency (how near-optimal will the solution be and how likely is this to happen) and cost (in terms of the time that the process will take). A global optimum will not in general be possible with probability 1. In the case of SA, only at zero temperature the system, if in equilibrium, is in some absolute optimum with probability 1, but this equilibrium requires, to be reached, a very slow cooling law (see section 2, Eq. (4)) which turns, again, the algorithm inefficient.

More precisely, the condition to be satisfied by the transition probability matrix P is that if

$$P(I|J)\mathbf{j}(j) = P(J|I)\mathbf{j}(I)$$

being I and J any two elements in the state space and \mathbf{j} a function of a single state, then

$$\mathbf{r}(I,n) \xrightarrow{n \rightarrow \infty} \mathbf{j}(I) \quad \forall I \quad (6)$$

where $\mathbf{r}(I,n)$ is the probability of being in state I at time n .

This condition is called the *detailed balance* condition [1]. It is fulfilled by the SA due essentially to the ergodicity of its transition probability (section 2, condition iii). In this case, we have $\mathbf{j}(I) = \exp(F(I)/kT)$ for all I ; then the process leads \mathbf{r} to the equilibrium distribution of the Gibbs canonical ensemble.

Hence the proposed line of research, in its more technical concern, implies to find EA's with genetic operators that define a transition probability function satisfying detailed balance. For this to be possible, therefore, the genetic operators must be ergodic (i.e. without any transient subsets). Then, knowing the corresponding transition probability, it is possible to compute the asymptotic probability distribution of the algorithm.

4. CONCLUSIONS

We have presented a thermodynamic approach to Evolutionary Algorithms, on the base of the analysis of some features of the Simulated Annealing algorithm. This optimisation technique, despite its different origin and inspiration, can be considered as a simple example of evolution strategy and, hence, used as a case of study.

Provided it is known that no EA can be guaranteed to converge to a globally optimal population, unless infinite population is assumed, we claim that the question whether can EA's be guaranteed to

converge to a globally optimal population with a certain predetermined probability (i.e. to know a priori the asymptotic probability distribution) is not a purely theoretical one, since it may substantially affect their efficiency both in the accuracy of the solution obtained for a problem and in the time spent in getting that solution. The possibility to know a priori the probability of global convergence and, moreover, the ability to control it via a proper parameter setting is a quite attractive perspective, especially for those applications in which an optimal or near-optimal solution is needed.

We summarised some main results concerning equilibrium distribution of the SA (i.e. at fixed finite temperature) as well as its asymptotic distribution for the case in which equilibrium is reached at zero temperature. In any case, the equilibrium probability distribution is the Gibbs distribution. The temperature acts as a parameter: the lower it is, the higher the probability for the system to be near an optimal configuration once equilibrium has been reached. Analysing the conditions that make this behaviour possible, we find that, from the point of view of thermodynamics of equilibrium, the essential condition is detailed balance and, in its turn, this condition depends, from a dynamical systems point of view, on the ergodicity intrinsic to the transition probability matrix associated with the SA algorithm.

Therefore, a way for the design of EA's with such good properties (asymptotic convergence to maxima of the fitness function) could be to find genetic operators that define a transition probability function satisfying the condition of detailed balance. For this to be possible, in its turn, these operators must be ergodic. Thus, knowing the corresponding transition probability, it is possible to compute the asymptotic probability distribution of the algorithm.

It is worth to mention that many attempts have been made so far in order to design more sophisticated SA-based evolution strategies. The idea of a parallel SA has been dealt with by several authors. There are, however, several problems to be solved, the rate of convergence not being the least. In fact, some well known theoretical results [4] assert that for a parallel SA with p processors, the expected speedup in the convergence is of order $O(\log p)$. Therefore, even in the theoretically better founded case of SA, it is not clear how to efficiently integrate it into more complex evolutionary computing architectures. And for the same theoretical framework being applicable to more general EA's with large populations and complex genetic operators, the questions to be solved are even many more.

5. REFERENCES

- [1] Amit, D. J.: Modeling Brain Function. Cambridge, Cambridge University Press (1992).

- [2] Baeck, T., Fogel, D. B. and Michalewicz, Z.: Evolutionary Computation, two vols., London, Institute of Physics (2000).
- [3] Baeck, T., Hoffmeister, F. and Schwefel, H.-P.: A Survey of Evolution Strategies. Proc. Fourth Int. Conf. Genetic Algorithms, San Mateo, Morgan Kaufmann Publishers (1991) 2-9.
- [4] Baeck, T. and Schwefel, H.-P.: An overview of Evolutionary Algorithms for parameter optimization. Evolutionary Computation **1**(1), (1993) 1-23.
- [5] Davis, R. and Principe: A Markov Chain Analysis of the Genetic Algorithm. Evolutionary Computation **1**(3) (1993) 269-288.
- [6] Eiben, A. E., Smith, J. E., Eiben, Agoston E. and Smith, J. D.: Introduction to Evolutionary Computing, Springer (2003).
- [7] Goertzel, Ben. Chaotic Logic: Language, Thought and Reality from the Perspective of Complex Systems Science. New York: Plenum (1994).
- [8] Michalewicz, Z.: Genetic Algorithms + Data Structures = Evolution Programs. 3rd edn. Springer Verlag, Berlin Heidelberg New York (1996).
- [9] Mitra, D., F. Romeo and A. Sangiovanni-Vincentelli: Convergence and FiniteTime Behavior of Simulated Annealing. Adv. Appl. Prob. **18** (1986) 747-771.
- [10] Rechenberg, I.: Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution. Stuttgart, Frommann-Holzboog Verlag (1973)
- [11] Salamon, P., Sibani, P. and Frost, R.: Facts, Conjectures, and Improvements for Simulated Annealing, Siam Monographs on Mathematical Modeling and Computation (2002).
- [12] Schwefel, H.-P.: Numerische Optimierung von Computer-Modellen mittels der Evolutionsstrategie. Interdisciplinary Systems Research; 26. Basel: Birkhauser (1977).