A Comparative Study of Probability Collectives Based Multi-agent Systems and Genetic Algorithms

We compare Genetic Algorithms (GA's) with Probability Collectives (PC), a new framework for distributed optimization and control. In contrast to GA's, PC-based methods do not update populations of solutions. Instead they update an explicitly parameterized probability distribution p over the space of solutions. That updating of p arises as the optimization of a functional of p. The functional is chosen so that any p that optimizes it should be p peaked about good solutions. The PC approach works in both continuous and discrete problems. It does not suffer from the resolution limitation of the finite bit length encoding of parameters into GA alleles. It also has deep connections with both game theory and statistical physics. We review the PC approach using its motivation as the information theoretic formulation of bounded rationality for multi-agent systems. It is then compared with GA's on a diverse set of problems. To handle high dimensional surfaces, in the PC method investigated here p is restricted to a product distribution. Each distribution in that product is controlled by a separate agent. The test functions were selected for their difficulty using either traditional gradient descent or genetic algorithms. On those functions the PC-based approach significantly outperforms traditional GA's in both rate of descent, trapping in false minima, and long term optimization.

Typically the search of adaptive, distributed agent-based algorithms is conducted by having each agent run its own reinforcement learning algorithm [1]. In this methodology the global utility function $G(x)$ in the system maps a joint move of the agents, $x \in X$, to the performance of the overall system. However, in practice the agents in a MAS are bounded rational; the equilibrium they reach typically involves mixed strategies rather than pure strategies – i.e., they don't settle on a single point x optimizing $G(x)$. This suggests formulating a framework to explicitly account for the bounded rational, mixed strategy character of the agents. Probability Collectives (PC) adopts this perspective to show that the equilibrium of a MAS is the minimizer of a Lagrangian $\mathcal{L}(P)$ (derived using information theory) that quantifies the expected value of G for the joint distribution $P(x_1, x_2, ..., x_N)$.

Now consider a bounded rational game in which the agents are independent, with each agent i choosing its move x_i at any instant by sampling its probability distribution (mixed strategy) at that instant, $q_i(x_i)$. Accordingly, the probability distribution of the joint-moves is a *product distribution*; i.e., $P(x) = P(x_1, x_2, ..., x_N) = \prod_{i=1}^{N} q_i(x_i)$, if there are N agents participate in the game. In this representation of a MAS, lacking the full joint probability distribution, all coupling between the agents occurs indirectly. It is the separate distributions of the agents $\{q_i\}$ that are statistically coupled, while the actual moves of the agents are independent.

The core of PC-based algorithms is thus to approximate the joint distribution by the product distribution, and to concentrate on how the agents update the probability distributions across their possible actions instead of specifically on the joint action generated by sampling those distributions.

The PC approach differs from traditional optimization methods such as gradient descent or GA which concentrate on a specific choice for the design variables (i.e. pure strategies) and on how to update that choice. Since the PC approach operates directly on probability distributions, it offers a direct treatment for incorporating uncertainty, which is also represented through probabilities [2]. This is the most salient feature that this class of algorithms possesses – the search course is guided by a probability distribution over x, rather than a single value of x. By building such a probabilistic model of promising solutions and sampling the built model to generate new candidate solutions, PC allows the agents to significantly expand the range of exploration of the search space, and simultaneously focus on promising solutions areas. As a result, the estimation of distribution algorithms can provide a robust and scalable solution to many important classes of optimization problems. The following sections report a comparative study of the PC-based MAS and GA using several test problems.

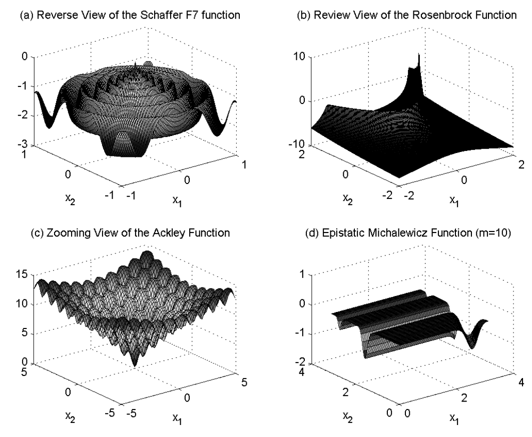

Figure 1: Surface plot for the four testbeds.

The first test function is Schaffer's test function F_7, which is defined as:

$$f(\overline{x}) = (x_1^2 + x_2^2)^{0.25}[sin^2(50(x_1^2 + x_2^2)^{0.1}) + 1],$$

where $-1 \leq x_i \leq 1$ for $1 \leq i \leq 2$. Figure 1.a displays the surface which is plotted upside down for easier viewing of the inverted minimum as a peak.

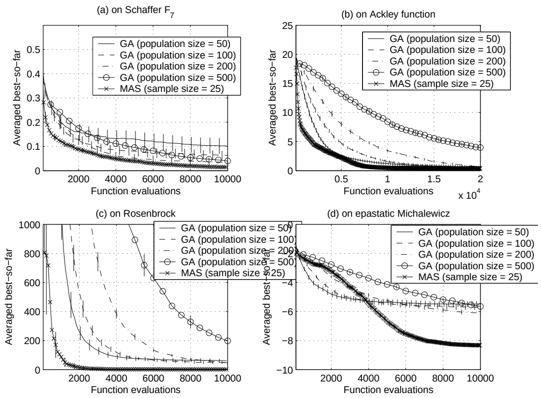

Figure 2: Best-so-far performance comparison.

Figure 2.a displays the best-so-far values attained by the Multi-agents system(MAS) and the GA as a function of the number of sample evaluations of the objective function. The methods distinguish themselves with different rates of initial descent of the objective function (on left) and the long-term performance (on right). Notably, the run-to-run variation of the performance trajectory is much lower on the PC-based MAS than for the GA (see vertical bars).

The second testbed is Ackley's Path, which is a widely used multimodal test function. The function's definition is:

$$f(\overline{x}) = -ae^{-b(\frac{\sum_i x_i^2}{N})^{\frac{1}{2}}} - e^{\frac{\sum_i cos(cx_i)}{N}} + a + e^1,$$

where a=20, b=0.2, $c = 2\pi$, and $-32.768 \leq x_i \leq 32.768$ for $1 \leq i \leq n$.

Figure 1.c gives a visual gist of the function in a lower 2-dimensional form. The surface is overall a single deep well with a locally rough surface. The empirical results of the search algorithms on this surface are displayed in Figure 2.b. It is clear that the PC-based MAS technique again significantly outperforms the GA in early decent towards the minimum.

The third testbed is the generalized Rosenbrock function in ten dimensions. The definition of this function is:

$$f(\overline{x}) = \sum_{i=1}^{N-1}[100(x_i^2 - x_{i+1})^2 + (1 - x_i)^2],$$

where $\overline{x} = [x_1, x_2, \ldots, x_N]^T$, $-5.12 \leq x_i \leq 5.12$.

Rosenbrock's saddle is a classic optimization problem with a narrow global optimum hidden inside a long, narrow, curved flat valley. Monte-carlo methods will have difficulty landing a point in the narrow spike and thus will not efficiently locate it. The U-shape will also tend to make decomposition of the PC into a product distribution challenging. Since it has no barriers the surface would be ripe for gradient descent; however while the valley will be found quickly the curvature and flatness of the valley floor will frustrate sampled gradient estimation. The empirical results on this surface are displayed in Figure 2.c. It is again clear that the PC-based MAS technique again significantly outperforms the GA in early decent towards the minimum.

The final testbed employed in this section is Michalewicz's epistatic function:

$$f(\overline{x}) = -\sum_{i=1}^{N} sin(y_i)sin^{2m}(\frac{iy_i^2}{\pi}),$$

where

$$\begin{aligned} y_i &= x_i cos\frac{\pi}{6} - x_{i+1} sin\frac{\pi}{6}, \quad \text{if } i \bmod 2 = 1 \text{ and } i \neq N; \\ y_i &= x_{i-1} sin\frac{\pi}{6} + x_i cos\frac{\pi}{6}, \quad \text{if } i \bmod 2 = 0 \text{ and } i \neq N; \\ y_N &= x_N, \end{aligned}$$

$0 \leq x_i \leq \pi$ for $1 \leq i \leq N$.

A system is highly epistatic if the optimal allele for any locus depends on a large number of alleles at other loci. This function is a highly multimodal, nonlinear and nonseparable testbed ($n!$ local optima). A sketch of a two-dimensional version of this function is displayed in Figure 1.d for the steepness parameter $m = 10$. Larger m leads to more difficult search. For very large m the function behaves like a needle in the haystack since the function values for points in the space outside the narrow peaks give very little information on the location of the global optimum.

The empirical results on this surface are displayed in Figure 2.d (for $m = 200$). It is clear that the PC-based MAS technique again significantly outperforms the GA. In particular, the PC still demonstrates a surprising search power even though the function behaves like a needle in the haystack and is very difficult to search.

References

[1] D. H. Wolpert and K. Wheeler and K. Tumer (2000), *Europhysics Letters*, **49 (6)**.

[2] Bieniawski, S. and Wolpert, D. H. and Kroo, I. 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, AIAA Paper 2004-4580.

Contact Information: Chien-Feng Huang – Center for Nonlinear Studies, Los Alamos National Laboratory, MS-B258, Los Alamos, NM, 87545, Phone: (505) 665-9176, email: cfhuang@lanl.gov.