

# Expected improvements for the asynchronous parallel global optimization of expensive functions : potentials and challenges

J. Janusevskis<sup>1</sup>, R. Le Riche<sup>1,2</sup>, D. Ginsbourger<sup>3</sup>, and R. Girdziusas<sup>1</sup>

<sup>1</sup> Ecole des Mines de Saint-Etienne, Saint-Etienne, France,

<sup>2</sup> CNRS, UMR 5146 Cl. Goux, France,

<sup>3</sup> Departement of Mathematics and Statistics, University of Bern, Switzerland.

**Abstract.** Sequential sampling strategies based on Gaussian processes are now widely used for the optimization of problems involving costly simulations. But Gaussian processes can also generate parallel optimization strategies. We focus here on a new, parameter free, parallel expected improvement criterion for asynchronous optimization. An estimation of the criterion, which mixes Monte Carlo sampling and analytical bounds, is proposed. Logarithmic speed-ups are measured on 1 and 9 dimensional functions.

## 1 Introduction to parallel expected improvements

The current technological solution for optimizing functions of numerically costly simulators is to rely on an increasing number of processing units (processors, cores, GPUs). Demanded features of new parallel optimization algorithms are not only high speed-ups but also the ability to work with heterogeneous processing units (e.g., computing grids) and fault tolerance. Evolutionary algorithms offer many opportunities for parallelization, including in heterogeneous computing networks, and in master-worker or island computing structures [1]. Master-worker parallel optimization algorithms are more common as they fit the costly simulator case: the optimizer is the master node, the workers evaluate the objective functions and constraints. In [5] for example, a parallel and asynchronous version of the local pattern search algorithm is described. [8] presents a deterministic global parallel algorithm which, contrarily to the forthcoming method is based on radial basis functions and has a synchronization step.

This article proposes a sampling criterion based on Gaussian processes. It is an asynchronous extension of the multi-points expected improvement of [2] which, itself, was a parallel extension of the expected improvement of [7]. Let us consider the optimization problem  $\min_{x \in \mathbb{R}^d} f(x)$  where each evaluation of  $f$  implies a call to a numerically intensive simulation program, and assume that a set of past observations  $(\mathbb{X}, f(\mathbb{X}))$  has been gathered. In the Bayesian Global Optimization settings considered here, the unknown  $f$  is represented *a priori* by a Gaussian Process  $(Y(x))_{x \in \mathbb{R}^d}$ , and is being approximated relying on the conditional distribution of  $Y$  knowing that  $Y(\mathbb{X}) = f(\mathbb{X})$  (*Kriging metamodel*).

We focus here on cases where  $\lambda$  computing nodes are available for starting new simulations while  $\mu$  computing nodes are currently evaluating  $f$  at a set of  $\mu \geq 0$  “busy” points  $\mathbb{X}_{\text{busy}} := \{x_b^1, \dots, x_b^\mu\}$ . We define the **asynchronous multi-points Expected Improvement** of  $\mathbb{X}_{\text{asy}} := \{x_a^1, \dots, x_a^\lambda\}$  as

$$EI^{(\mu, \lambda)}(\mathbb{X}_{\text{asy}}) := \mathbb{E} \left[ (\min(Y(\mathbb{X} \cup \mathbb{X}_{\text{busy}})) - \min(Y(\mathbb{X}_{\text{asy}})))^+ | Y(\mathbb{X}) = f(\mathbb{X}) \right],$$

where  $[\bullet]^+ \equiv \max(0, \bullet)$ . This criterion, which was initially introduced in [3], measures how much progress with respect to already calculated or currently calculating points will be made (on average) at  $\mathbb{X}_{\text{asy}}$ . The case  $\mu = 0, \lambda = 1$  is the usual  $EI$  defined in the EGO algorithm [7]. A parallel algorithm that works by maximizing  $EI^{(\mu, \lambda)}$  can now be introduced:

**Asynchronous Parallel  $EI$  algorithm**

1. Generate  $\mathbb{X}$  through a space-filling design. Calculate  $f(\mathbb{X})$ .
2. While calculation budget not exhausted do
  - (a) [non blocking] Retrieve new  $x$ ’s and  $f(x)$ ’s if any. Update the kriging model (optionally with parameter re-estimation). Update  $\lambda$  and  $\mu$ .
  - (b) Generate  $\lambda$  points by  $\max EI^{(\mu, \lambda)}(x_a^1, \dots, x_a^\lambda)$  using a global optimizer (e.g. CMA-ES [4]). Send them to worker nodes for evaluation.

The classic  $EI$  criterion has the desirable property that its maximum lies away from already sampled points of  $\mathbb{X}$  and strikes a compromise between the exploration of unknown regions of the design space and the intensification of the search in known highly performing regions [7]. In addition, the multi-points asynchronous  $EI^{(\mu, \lambda)}$  has its maximum away from any subset of already sampled and currently running points,  $\mathbb{X} \cup \mathbb{X}_{\text{busy}}$  (it is null there while  $I$  is a positive variable, [6]). Another advantage of  $EI^{(\mu, \lambda)}$  is that it does not introduce extra parameters.

$EI^{(\mu, \lambda)}$  accounts for all ratios of optimizer over simulation computation times. If the simulations are much longer than any optimizer iteration (kriging update and  $EI$  maximization), points will be allocated to newly available nodes one at a time, in which case  $EI^{(\mu, \lambda=1)}$  will be used. Vice versa, if the optimizer iteration takes longer than the simulations, no busy point occurs and  $EI^{(\mu=0, \lambda)}$  is relevant. Intermediate cases call for general  $EI^{(\mu, \lambda)}$ ’s.

## 2 Bounds and estimation of $EI^{(\mu, \lambda)}$

While the calculation of  $EI^{(0, 1)}$  and  $EI^{(0, 2)}$  is analytical [2], no general expression for  $EI^{(\mu, \lambda)}$  is known. We propose to estimate  $EI^{(\mu, \lambda)}$  through a Monte Carlo strategy augmented by bounds knowledge. In particular, the following bounds are established in [6] where  $EI^*(x_b^i, x_a^j) := \mathbb{E} [(Y(x_b^i) - Y(x_a^j))^+ | Y(\mathbb{X}) = f(\mathbb{X})]$  :

$$0 \leq EI^{(\mu, \lambda)}(\mathbb{X}_{\text{asy}}) \leq \min \left( \sum_{i=j}^{\lambda} EI(x_a^j), \sum_{j=1}^{\lambda} EI^*(x_b^1, x_a^j), \dots, \sum_{j=1}^{\lambda} EI^*(x_b^\mu, x_a^j) \right)$$

All the expressions in the bounds are analytical, including the  $EI^*$  terms, because they are instances of the usual  $EI$  formula [7].

Now, a straightforward way of getting approximations of  $EI^{(\mu,\lambda)}(\mathbb{X}_{\text{asy}})$  consists in relying on Monte Carlo estimates  $EI_{MC}^{(\mu,\lambda)}(\mathbb{X}_{\text{asy}})$  obtained by conditional simulations of the process  $Y$ . A simple reasoning on Bayes law proves that  $EI^{(\mu,\lambda)}(\mathbb{X}_{\text{asy}})|EI_{MC}^{(\mu,\lambda)}(\mathbb{X}_{\text{asy}})$  can be consistently approximated by a truncated Gaussian distribution, with mean  $M(\mathbf{x})$  and variance  $V^2(\mathbf{x})$  given in [6]. The asynchronous parallel EI algorithm introduced earlier has a step that maximizes  $EI^{(\mu,\lambda)}(\bullet)$  with the CMA-ES algorithm ([4]). In CMA-ES, the objective function values are used to rank the explored points. We propose to modify this comparison of points (say  $\mathbf{x}^i$  and  $\mathbf{x}^j$ ) in order to control the Monte Carlo simulations:

**Pairwise ranking procedure**

1. Set confidence  $k$  (e.g., over 60% confidence for  $k = 1$ ),  $N$ ,  $\Delta N$ ,
2. While not stop do
  - (a) If  $M(\mathbf{x}^i) - kV(\mathbf{x}^i) \geq M(\mathbf{x}^j) + kV(\mathbf{x}^j)$  then  $EI^{(\mu,\lambda)}(\mathbf{x}^i) \geq EI^{(\mu,\lambda)}(\mathbf{x}^j)$ , stop. (And vice versa.)
  - (b) Decrease the Monte Carlo variances by allocating extra samples:  $\Delta N_i = V(\mathbf{x}^i)\Delta N/(V(\mathbf{x}^i) + V(\mathbf{x}^j))$ ,  $\Delta N_j = \Delta N - \Delta N_i$ .

### 3 Test results




Fig.1: Left: two of the ten sampled trajectories with Matern 5/2 kernel and scale  $\theta = 0.15$ . Right: mean normalized improvement as a function of iteration for  $\lambda = 1$  to 4 ( $\mu = 0$ ), averaged over 10 test functions (defined as 1D trajectories such as those on the left).

The tests reported in [6] are now summarized.

Firstly, 100 functions in 1D have been generated by sampling Gaussian process trajectories (see the 2 examples of Fig.1). For each function, a pair of points was randomly chosen and their  $EI^{(\mu,\lambda)}$  compared,  $\mu = 0, 1, 3$  and  $\lambda = 1$  to 5. It is observed that, on average, accounting for the bounds divides by 7 the total number of Monte Carlo simulations necessary to discriminate the points. However, a first difficulty appeared. When  $\mu > 0$ , 25% of the pairs needed over  $N = 100,000$  MC samples for the comparison, i.e., their points had very close  $EI^{(\mu,\lambda)}$  values: the  $EI^{(\mu,\lambda)}$  function has plateaus. This proportion increased when the pairs were generated by optimization because these plateaus correspond to high performance regions of the design space.

The second issue stems directly from  $EI^{(\mu,\lambda)}$ 's definition: its input is high dimensional ( $\dim \mathbb{X}_{\text{asy}} = \lambda d$ ), making its maximization potentially costly.

Thirdly, on the positive side, the experiments summarized in Fig.1 and 2 show that logarithmic speed-ups are obtained: we have observed  $\frac{\text{time to solve, } \lambda=1}{\text{time to solve, } \lambda} \approx 1 + b \log(\lambda)$  where  $b = 0.85$  and  $0.49$  in 1 and 9 dimensions, respectively. The 9D test function is the approximation of a matrix by its rank 1 counter part:

$$(\mathbf{B}_{mr}^*, \mathbf{C}_{rn}^*) = \arg \min_{\mathbf{B}_{mr}, \mathbf{C}_{rn}} \|\mathbf{A}_{mn} - \mathbf{B}_{mr} \mathbf{C}_{rn}\|, \quad r < \text{rank}(\mathbf{A}_{mn}), \quad (1)$$

where  $\|\cdot\|$  is the Frobenius norm,  $\mathbf{A}$  is a  $4 \times 5$  matrix of uniformly distributed elements in  $[0, 1]$ ,  $\mathbf{B}$  and  $\mathbf{C}$  are a column and a row vector whose elements are constrained to be in  $[-1, 1]$ , respectively. This is a continuous nonconvex 9-dimensional box-constrained optimization problem whose solution is given by the first singular vectors. An instance of a typical run is shown for  $\lambda = 1$  to 16,  $\mu = 0$ , in Fig. 2.

This speed-up is due to the kriging model which summarizes all of the information gathered by the worker nodes, including the currently running simulations. Further work addressing the aforementioned estimation and optimization difficulties of  $EI^{(\mu,\lambda)}$  is needed.

## References

1. Branke, J., Kamper, A., Schmeck, H.: Distribution of evolutionary algorithms in heterogeneous networks. In Deb K. et al., eds, GECCO 2004, LNCS 3102, Springer, pp. 923-934, 2004.
2. Ginsbourger, D., Le Riche, R., Carraro, L.: Kriging is well-suited to parallelize optimization. In Tenne Y. and Goh C.-K., eds., Computational Intelligence in Expensive Optimization Problems, Springer series in Evolutionary Learning and Optimization, pp. 131-162, 2009.
3. Ginsbourger, D., Janusevskis, J., Le Riche, R.: Dealing with asynchronicity in parallel Gaussian Process based global optimization. HAL technical report no. hal-00507632, <http://hal.archives-ouvertes.fr/hal-00507632/>, July 2010
4. Hansen, N., Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies. Evolutionary Computation, 9(2), pp. 159-195, 2001.
5. Kolda, T. G. : Revisiting asynchronous parallel pattern search for nonlinear optimization. SIAM J. Optimization, 16(2), pp. 563-586, 2005.




Fig. 2: Mean normalized improvement as a function of iteration for  $\lambda = 1$  to 16 ( $\mu = 0$ ), 9 dimensional test case.

6. Janusevskis, J., Le Riche R., Ginsbourger, D.: Parallel expected improvements for global optimization: summary, bounds and speed-up. HAL technical report no. hal-00613971, [http://hal.archives-ouvertes.fr/hal-00613971\\_v1](http://hal.archives-ouvertes.fr/hal-00613971_v1), August 2011
7. Jones, D.R., Schonlau, M. and Welch, W. J.: Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4), pp. 455-492, 1998.
8. Regis, R. G., Shoemaker, C. A.: Parallel radial basis function methods for the global optimization of expensive functions. *European J. of Operational Research*, 182, pp. 514-535, 2007.