Global extremum seeking by Kriging with a multi-agent system

Arthur Kahn * **Julien Marzat** * **Hélène Piet-Lahanier** *
Michel Kieffer **

* *ONERA - The French Aerospace Lab, F-91123 Palaiseau, France, firstname.lastname@onera.fr*
** *L2S, CNRS-SUPELEC-Univ-Paris-Sud, F-91192, France, firstname.lastname@lss.supelec.fr*

Abstract This paper presents a method for finding the global maximum of a spatially varying field using a multi-agent system. A surrogate model of the field is determined via Kriging (Gaussian process regression) from a set of measurements collected by the agents. A criterion exploiting Kriging statistical properties is introduced for selecting new sampling points that each vehicle must rally. These new points are obtained as a compromise between improvement of the estimate of the global maximum and traveling distance. A cooperative control law is proposed to move the agents to the desired sampling points while avoiding collisions. Simulation results show the interest of the method and how it compares with a state-of-the-art solution.

Keywords: Kriging, global extremum seeking, multi-agent systems, cooperative exploration.

1. INTRODUCTION

Recent years have witnessed a growing development of autonomous multi-agent systems (MAS) such as unmanned aerial vehicles (UAV) or unmanned ground vehicles (UGV). The main advantage of MAS is that agents can cooperate to fulfill some mission more efficiently than a single agent (Choi et al. (2009)). Typical missions are field exploration or maximum seeking (Williams and Sukhatme (2012)). A gradient climbing approach is often employed to bring the vehicles to the extremum (Choi and Horowitz (2007); Ogren et al. (2004)), based on a cooperative estimation over the MAS (Zhang and Leonard (2010)). The main drawback of gradient climbing is its local convergence. As the maximum seeking is performed using a MAS, one must also take into account limitations due to the vehicle dynamics since the agents have to move to some desired position to measure the field value.

The aim of this paper is to search for a global optimum of an unknown field with a MAS, based on a Kriging model (or Gaussian process regression) computed from sampling points (Jones (2001)), while taking into account the individual agent dynamics. Kriging is an interpolation method that gives the best linear unbiased prediction with prior assumption on the covariance of the estimated field.

Few papers have used a Kriging model in the context of MAS. In Cortes (2009), Kriging is used to model a time-varying field and estimate its gradient to drive the agents to the field maximum. Gu and Hu (2012) partition the space with Voronoi cells based on a Kriging model of the field. The agents are controlled to minimize the uncertainty of the model without explicitly searching for the field maximum. In Karasev (2012), Gaussian process regression is used to define the positions of the agents to improve visual search. In Choi et al. (2008), the agents stay in formation and several criteria based on the Kriging estimate are presented to perform field exploration or maximum seeking. At each instant the criterion is updated and a new direction is considered for the agents. The authors improved the method in Xu et al. (2011) for field exploration. Each vehicle tries to minimize the cumulated estimated field uncertainty evaluated on a grid of target points. Centralized and distributed schemes were presented and compared.

The proposed approach performs a trade-off between field exploration where uncertainty is high and global maximum seeking. Measurement effort is directed to areas that can potentially improve the current estimate of the global maximum, unlike other methods that focus on exploration (which is time-consuming) or gradient climbing (which may find only a local maximum). The Kriging model is well suited to determine whether an area must be explored, thanks to the estimation uncertainty model it provides.

This paper is organized as follows. Section 2 formulates the problem. Section 3 presents the Kriging estimator and related optimization criteria. Section 4 defines the proposed criterion and control law for field exploration and maximum seeking. Simulation results are described in Section 5 to compare the proposed approach with that of Xu et al. (2011).

2. PROBLEM FORMULATION

Consider some unknown, continuous, and time-invariant scalar field ϕ defined over a compact space $D \subset \mathbb{R}^m$, which has to be maximized. The field ϕ may present several local extrema but is assumed to have a unique argument of its global maximum.

The aim is to perform the maximization of ϕ with a fleet of N identical mobile agents. Each agent $i \in 1 \dots N$ is able to measure ϕ and its current location $\mathbf{x}_i(t)$ at any time instant t. The search for the maximum is operated by moving the agents in D. The dynamics of each agent is modeled as

$$M\ddot{\mathbf{x}}_i + C(\mathbf{x}_i, \dot{\mathbf{x}}_i)\dot{\mathbf{x}}_i = \mathbf{u}_i, \quad (1)$$

where M is the mass of the agent, $C(\mathbf{x}_i, \dot{\mathbf{x}}_i)$ a non negative friction coefficient (Wang (1991)), and $\mathbf{u}_i(t)$ their control input. The field measurement model is

$$y_i = y(\mathbf{x}_i) = \phi(\mathbf{x}_i) + w_i, \quad (2)$$

where w_i is a zero-mean Gaussian noise with known variance σ_w^2.

The agents are synchronized and time is discretized with a period T. At each $t_k = kT$, they compute their own control input and may collect a field measurement and transmit it, together with their location. Two agents are assumed to be able to communicate when their distance is smaller than r. The set of neighbors of agent i at time t_k is denoted

$$\mathcal{N}_i(t_k) = \{j \mid \|\mathbf{x}_i - \mathbf{x}_j\| \leqslant r\}. \quad (3)$$

This defines an undirected and time-varying communication graph. Communications are assumed to be lossless and without delay. At time t_k, each agent i possesses a set $S_i(t_k)$ gathering its own past field measurements and those received from its neighbors. Let $\mathcal{M}(t_k)$ be the set of the agents that collect a measurement at time t_k. The set $S_i(t_k)$ is then defined as

$$S_i(t_k) = \bigcup_{l=0}^{k} \{[y_j(t_l), \mathbf{x}_j(t_l)] \mid j \in \mathcal{N}_i(t_l) \cap \mathcal{M}(t_l)\}. \quad (4)$$

Maximum seeking is performed with the help of a model of ϕ taking the form of a Gaussian process, whose parameters are estimated from $S_i(t_k)$. The model has to provide simultaneously a mean estimate $\hat{\phi}_{i,k}$ of the field over the space D and an estimation variance $\sigma^2_{\phi,i,k}$. Each agent $i \in [1, \ldots, N]$ updates its model of the field at time t_k only if $S_i(t_{k-1}) \neq S_i(t_k)$.

The mission goal is to find the location $\mathbf{x}_M$ of the maximum of ϕ on D,

$$\mathbf{x}_M = \arg\max_{\mathbf{x}\in D} \{\phi(\mathbf{x})\}. \quad (5)$$

Since ϕ is unknown, the search is performed on $\hat{\phi}_{i,k}$, which is updated iteratively in such a way that, there exists $K_\epsilon > 0$ such that,

$$\hat{\mathbf{x}}_M^i(t_k) = \arg\max_{\mathbf{x}\in D} \left\{\hat{\phi}_{i,k}(\mathbf{x})\right\} \quad (6)$$

satisfies $||\hat{\mathbf{x}}_M^i(t_k) - \mathbf{x}_M|| < \epsilon$ for all $k \geqslant K_\epsilon$. For this purpose, $\hat{\phi}_{i,k}(\mathbf{x})$ has to be updated by properly selecting field measurement points for each agent. Basics of Kriging and various related policies to select such points are presented in Section 3, before introducing the proposed method in Section 4.

3. KRIGING

3.1 Basics of Kriging

Consider the scalar function

$$\phi : \mathbf{x} \in D \rightarrow \phi(\mathbf{x}) \in \mathbb{R}. \quad (7)$$

ϕ is modeled by the Gaussian process

$$\Omega(\mathbf{x}) = \mathbf{r}(\mathbf{x})^T\boldsymbol{\beta} + Z(\mathbf{x}), \quad (8)$$

where $\mathbf{r}$ is some regression vector and $\boldsymbol{\beta}$ is a parameter vector (a linear regression model is used in this paper). Z is a zero-mean Gaussian process with covariance function modeled as

$$\text{cov}(Z(\mathbf{x}), Z(\mathbf{x}')) = \sigma_z^2 \xi(\mathbf{x}, \mathbf{x}'), \quad (9)$$

where

$$\xi(\mathbf{x}, \mathbf{x}') = \exp\left[-\sum_{i=1}^{m} \frac{1}{\theta_i} |\mathbf{x}(i) - \mathbf{x}'(i)|^{q_i}\right] \quad (10)$$

is a spatial correlation function (Schonlau (1997)). Its parameters are the nominal variance σ_z^2, the range of the spatial correlation θ_i, and $q_i \in [0,\ 2]$ which determines the smoothness of the interpolation. Those parameters depend on the characteristics of the field. They may be known *a priori* or have to be estimated.

Assume that some vector of n noisy measurements of ϕ

$$\mathbf{y} = [y(\mathbf{x}_1), \ldots, y(\mathbf{x}_n)]^T \quad (11)$$

has been collected, under measurement model (2).

Define vectors

$$\mathbf{R} = [\mathbf{r}(\mathbf{x}_1), ..., \mathbf{r}(\mathbf{x}_n)]^T, \quad (12)$$

$$\mathbf{k}_\mathbf{x} = [\xi(\mathbf{x}, \mathbf{x}_1), ..., \xi(\mathbf{x}, \mathbf{x}_n)]^T, \quad (13)$$

and the covariance matrix $\mathbf{K}$ with components

$$\mathbf{K}_{ij} = \xi(\mathbf{x}_i, \mathbf{x}_j) + \sigma_w^2 \mathbf{I}_n. \quad (14)$$

A linear estimator $\mathbf{a}_\mathbf{x}^T\mathbf{y}$ of $\Omega(\mathbf{x})$ is searched for. The bias of this estimator is

$$E\left[\Omega - \mathbf{a}_\mathbf{x}^T\mathbf{y}\right] = \mathbf{r}(\mathbf{x})^T\boldsymbol{\beta} - \mathbf{a}_\mathbf{x}^T\mathbf{R}\boldsymbol{\beta} \quad (15)$$

and its associated variance is

$$\begin{aligned} E[(\Omega(\mathbf{x}) - \mathbf{a}_\mathbf{x}^T\mathbf{y})^2] = (&\mathbf{a}_\mathbf{x}\mathbf{R}\boldsymbol{\beta} - \mathbf{r}(\mathbf{x})^T\boldsymbol{\beta}) \\ &+ \mathbf{a}_\mathbf{x}^T\sigma_z^2\mathbf{K}\mathbf{a}_\mathbf{x} + \sigma_z^2 - 2\mathbf{a}_\mathbf{x}^T\sigma_z^2\mathbf{k}_\mathbf{x} \end{aligned} \quad (16)$$

An unbiased estimator requires $\mathbf{a}_\mathbf{x}\mathbf{R}\boldsymbol{\beta} - \mathbf{r}(\mathbf{x})^T\boldsymbol{\beta} = 0$. The mean of the estimator is then

$$\hat{\phi}(\mathbf{x}) = \mathbf{r}(\mathbf{x})^T\boldsymbol{\beta} + \mathbf{k}_\mathbf{x}^T\mathbf{K}^{-1}(\mathbf{y} - \mathbf{R}\boldsymbol{\beta}) \quad (17)$$

and its associated variance :

$$\sigma_\phi^2(\mathbf{x}) = E[(\Omega(\mathbf{x}) - \mathbf{a}_\mathbf{x}^T\mathbf{y})^2] = \sigma_z^2(1 - \mathbf{k}_\mathbf{x}^T\mathbf{K}^{-1}\mathbf{k}_\mathbf{x}) \quad (18)$$

In the following, each agent i builds its own model from $S_i(t_k)$, resulting in estimated mean $\hat{\phi}_{i,k}(\mathbf{x})$ and variance $\sigma^2_{\phi,i,k}(\mathbf{x})$.

3.2 Sampling methods for global optimization

Different methods exist to choose where to select the next sampling point $\mathbf{x}^d$ based on Kriging for finding the global optimum of a field known only at sampled locations on a multivariate input space.

Kushner (1962) uses the Gaussian cumulative distribution function of the Kriging model to maximize the probability of improving the estimate of the maximum. This criterion promotes local extrema over exploration. Expected Improvement (Schonlau and Welch (1996)) exploits the Kriging model to determine the points which are the most likely to improve in average the global maximum estimate. To this end, it allows a trade-off between exploration and maximum search. The convergence to the global optimum of these strategies has been proven in Vazquez and Bect

(2010), providing that the correlation function satisfies some assumptions.

Alternatively, the lower confidence bounding function

$$\mathcal{C}_{\mathrm{lcb}}(\mathbf{x}) = \hat{\phi}(\mathbf{x}) + b\sigma_{\phi}(\mathbf{x}) \tag{19}$$

proposed in (Cox and John (1997)) is useful for finding a location

$$\mathbf{x}^d = \arg\max_{\mathbf{x}\in D} \mathcal{C}_{\mathrm{lcb}}(\mathbf{x}) \tag{20}$$

where ϕ may reach an extremum or presents a high level of uncertainty. Increasing the parameter b in (19) promotes points with higher uncertainty.

4. PROPOSED METHOD

4.1 Sampling point selection for MAS

The previous methods assume that a single next sampling point can be chosen arbitrarily in D. A specificity of sampling point selection for MAS comes from the need to determine a sampling point for *each* agent, which requires sharing the exploration load among agents. Moreover, sampling point selection has to account for the dynamics of each agent.

Assume that the estimate of the maximum of ϕ performed by agent i at time t_k is

$$f^i_{\max}(t_k) = \max_{\mathbf{x}\in S_i(t_k)} \{\hat{\phi}_{i,k}(\mathbf{x})\}. \tag{21}$$

The proposed technique, inspired from (19), consists in selecting the next sampling point for agent i as

$$\mathbf{x}^d_i(t_k) = \arg\min_{\mathbf{x}\in D} \left\{ J^{(k)}_i(\mathbf{x}) \right\} \tag{22a}$$

$$\text{s.t. } \hat{\phi}_{i,k}(\mathbf{x}) + b\sigma_{\phi,i,k}(\mathbf{x}) > f^i_{\max}(t_k) \tag{22b}$$

where

$$J^{(k)}_i(\mathbf{x}) = \|\mathbf{x}_i(t_k) - \mathbf{x}\|^2 - \sum_{j\in\mathcal{N}_i(t_k)} \alpha\|\mathbf{x}_j(t_k) - \mathbf{x}\|^2, \tag{23}$$

where α and b are two positive tuning parameters.

The constraint (22b) defines the subset of D potentially containing the global maximum. Using (23), one searches in this subset for a sampling point close to the current agent location $\mathbf{x}_i(t_k)$ and far enough from the other agent locations $\mathbf{x}_j(t_k)$, $j \in \mathcal{N}_i(t_k)$. However, the convergence to the maximum position is not guaranteed.

4.2 Control law

Once agent i has computed its next sampling point $\mathbf{x}^d_i(t_k)$, it should modify its trajectory to reach this point and collect a new measurement there.

For convenience and readability of the equations, time dependence in t_k is omitted in this section. The control law is designed so that each vehicle moves towards $\mathbf{x}^d_i$ with a desired velocity $\dot{\mathbf{x}}^d_i$, and avoids colliding with its neighbors. It is expressed as

$$\begin{aligned}\mathbf{u}_i =& C(\mathbf{x}_i,\dot{\mathbf{x}}_i)\dot{\mathbf{x}}_i - k_3(\mathbf{x}_i - \mathbf{x}^d_i) \\ &+ 2k_2\sum_{j=1}^{N}(\mathbf{x}_i - \mathbf{x}_j)\frac{1}{q}\exp\left(-\frac{(\mathbf{x}_i - \mathbf{x}_j)^T(\mathbf{x}_i - \mathbf{x}_j)}{q}\right) \\ &- k_1(\dot{\mathbf{x}}_i - \dot{\mathbf{x}}^d_i)\end{aligned} \tag{24}$$

with k_1, k_2 and k_3 three positive gains, q a parameter related to the safety distance between two vehicles of the fleet. For convenience,

$$g(\mathbf{x}_i,\mathbf{x}_j) = \exp\left(-\frac{(\mathbf{x}_i - \mathbf{x}_j)^T(\mathbf{x}_i - \mathbf{x}_j)}{q}\right) \tag{25}$$

is denoted g_{ij} and $\boldsymbol{\delta}_{ij} = \mathbf{x}_i - \mathbf{x}_j$ is the difference vector of position between vehicles i and j.

The stability of the control law is analyzed by considering

$$V = \frac{1}{2}\sum_{i=1}^{N}\left[\dot{\mathbf{x}}_i^T M\dot{\mathbf{x}}_i + (\mathbf{x}_i - \mathbf{x}^d_i)^T k_3(\mathbf{x}_i - \mathbf{x}^d_i) + k_2\sum_{j=1}^{N} g_{ij}\right] \tag{26}$$

which can be shown to be a Lyapunov function using derivations similar to that in (Cheah et al. (2009)).

Algorithm 1 summarizes the steps performed for maximum seeking. A measurement is only collected by agent i when it reaches its next sampling point $\mathbf{x}^d_i(t_k)$.

Algorithm 1 Maximum seeking Algorithm

```
for every time instant t_k do
    for each agent i do
        if x_i(t_k) = x_i^d(t_k) then
            Acquire a measurement y_i at x_i(t_k) (2)
            Broadcast {y_i, x_i(t_k)} to neighbors j ∈ N_i(t_k)
            Update S_i(t_k)
        end if
        if S_i(t_k) ≠ S_i(t_{k-1}) then
            Update the Kriging model (17) and (18)
            Solve (22) to find x_i^d(t_k)
        else
            x_i^d(t_k) = x_i^d(t_{k-1})
        end if
        Compute the control input u_i(t_k) (24)
    end for
end for
```

5. SIMULATION RESULTS

The proposed method is compared on a two-dimensional example, assuming a complete communication graph, to a reference technique described in Xu et al. (2011) and briefly summarized in what follows.

5.1 Reference method

The method proposed in Xu et al. (2011) evaluates the next sampling points for all agents in such a way that the sum of the variances of the Kriging model at some reference points is minimized. The next sampling points $\overline{\mathbf{x}} = [\mathbf{x}_1, ..., \mathbf{x}_N]^T$ are chosen by agent i to minimize

$$\mathcal{C}^i_{\mathrm{Xu}}(\overline{\mathbf{x}}) = \frac{1}{|\mathcal{J}|}\sum_{\boldsymbol{\xi}_j\in\mathcal{J}} \sigma^2_{\phi,i}(\overline{\mathbf{x}},\boldsymbol{\xi}_j) \tag{27}$$

where $\sigma^2_{\phi,i}(\overline{\mathbf{x}}, \boldsymbol{\xi}_j)$ is the predicted variance of the Kriging model of agent i at the target point $\boldsymbol{\xi}_j$ assuming that measurements are taken at $\overline{\mathbf{x}} = [\mathbf{x}_1, \ldots, \mathbf{x}_N]^T$. $\mathcal{J}$ is a set of properly chosen target points spread on D, and $|\mathcal{J}|$ is the cardinality of $\mathcal{J}$.

The target displacement speed $\dot{\mathbf{x}}_i^d$ of each agent is taken proportional to the gradient of (27), and achieved by using control law (24) with $\mathbf{x}_i^d(t_k) = \mathbf{x}_i(t_k)$. The agents collect measurements with a period equal to τ (two values are considered in the simulations). This is repeated iteratively until the criterion (27) falls below some threshold. At each period τ, the estimate of the maximum of the mean of the Kriging model can be computed.

5.2 Simulation conditions

In both cases, a fleet of $N = 3$ identical agents is considered, with $M = 1$ kg, $C = 0.001$ kg/s. The fleet is assumed to remain within a square field of size 50×50 m^2. Measurements are assumed noise free, $\sigma_w^2 = 0$. The communication graph is assumed to be fully connected, $r > 50\sqrt{2}$ m. The information is centralised as each agent has a complete knowledge of the others. The same Kriging model is used by all agents but depending on their current positions, the argument of the optimality criterion satisfying the constraints is different. A synthetic field ϕ_{sim} has been generated as the sum of three two-dimension Gaussian functions with maxima equal to 1.2, 1, and 1, located at [15 15], [40 35], and [10 35]. The covariance matrix for the three Gaussian functions is a diagonal matrix with values $\{100\ 100\}$. The global maximum of the field is equal to 1.25, located at [14.94 16.15].

The values of the parameters for the Kriging model are selected as $q_{1,2} = 2$, $\sigma_z = 0.5$, and $\theta_{1,2} = 50$. Linear regression was used for the deterministic part of the Gaussian process model. The sampling period is $T = 0.01$ s. The agents move using the control law defined in Section 4.2, with $q = 0.1$, $k_1 = 47$, $k_2 = 50$ and $k_3 = 1600$. The maximum velocity of the agents is fixed to 2 m/s.

The parameters used by the proposed approach are $b = 3$ in (22b) and $\alpha = \frac{1}{N} = \frac{1}{3}$ in (23). The optimization problem (22) is solved using the global optimizer DIRECT (Jones et al. (1993)). As the control law brings asymptotically the agents to the desired position $\mathbf{x}_i^d(t_k)$ with a null desired velocity $\dot{\mathbf{x}}_i^d(t_k) = 0$, it has been chosen to sample a measurement when $||\mathbf{x}_i(t_k) - \mathbf{x}_i^d(t_k)|| < \epsilon$ with $\epsilon = 0.01$ m.

For the reference method, two different values of the sampling period τ have been used: $\tau_1 = 5T$, and $\tau_2 = 20T$. There are 100 target points $\boldsymbol{\xi}_j$ in $\mathcal{J}$, uniformly distributed on a grid to cover the entire area D.

5.3 Simulation results

Figure 1 illustrates the evolution with time of the value of $G^k_{\max}(\mathbf{x}) = \hat{\phi}_k(\mathbf{x}) + b\sigma_{\phi,k}(\mathbf{x})$ for all $\mathbf{x} \in D$ with the proposed method and the synthetic field ϕ_{sim}. Red dots represent measurement locations. Figure 1(a) shows $G^k_{\max}$ after three measurements of each agent. In Figure 1(b), the agents spread to explore D. In Figure 1(c), the field

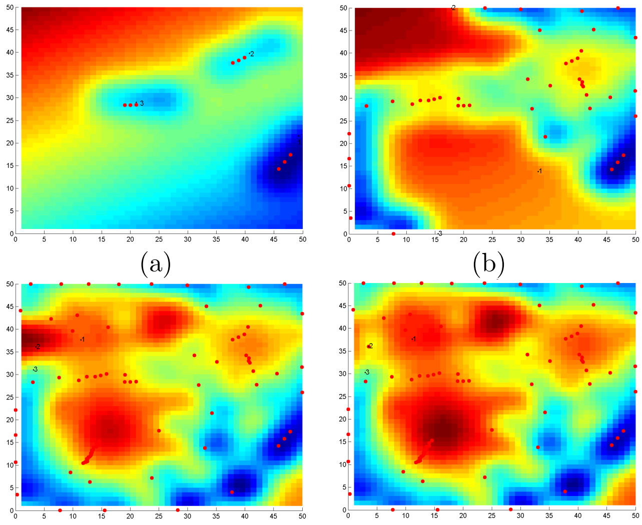

Figure 1. Maximum seeking by 3 agents with the proposed method

is mostly known and measurements have been performed near the true maximum location but the remaining level of uncertainty does not allow to end the mission. This requires the exploration of the few area still satisfying (22b). In Figure 1(d), no point satisfying (22b) can be found, so the search is stopped. Note that the variance (18) is not null everywhere: only areas of interest have been explored.

To compare the proposed method to the reference one, simulation results are averaged over several identical random initial locations of the agents in D. Figure 2 presents the evolution with time of the distance between the actual maximum location and the location of $f_{\max}(t_k)$. Figure 3 provides the number of measurements required by each method with respect to time. Figure 4 synthesizes the performances by showing the relation between the distance to the maximum and the number of measurements acquired. The results obtained with the proposed method are represented with blue curves, while the red and black ones are those obtained using the reference method respectively tuned with $\tau_1 = 5T$ and $\tau_2 = 20T$. Solid lines are averaged results while dotted lines correspond to the minimal and maximal values collected over all the runs.

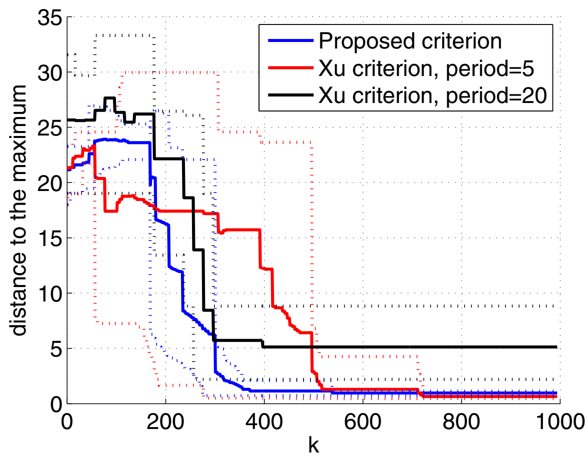

Figure 2. Distance to the maximum w.r.t. time

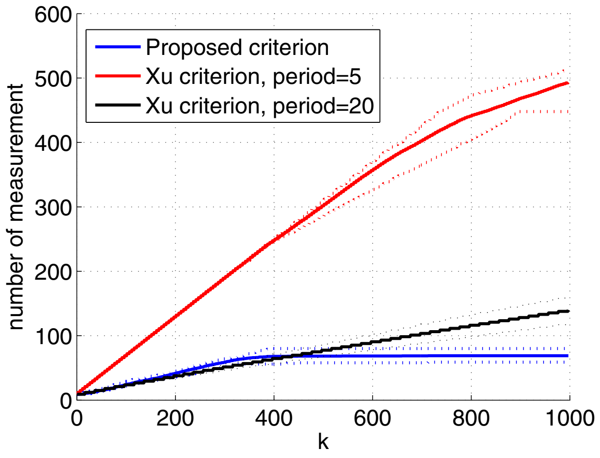

Figure 3. Number of measurements w.r.t. time

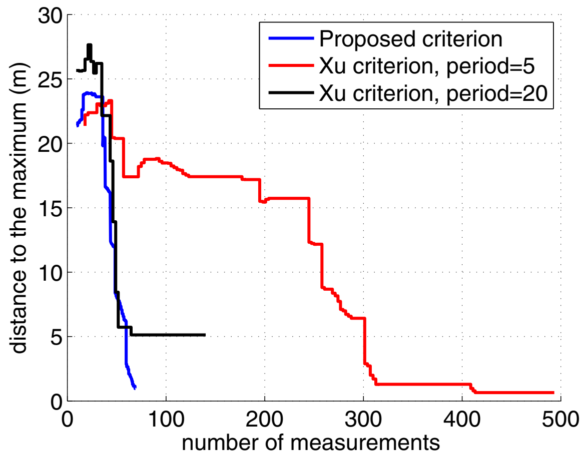

Figure 4. Distance to the maximum *vs* number of measurements

All methods present a similar dispersion of results. The reference method exhibits different characteristics depending on the choice of the sampling period τ. When τ is small, convergence to the maximum is precise but slow and the number of measurements is large. When τ is larger, the distance to the maximum decreases quickly but never converges, while few measurements are required. The proposed method does not need a tuning of τ and appears to combine all desired properties: a quick convergence to the maximum is achieved with few measurements.

While the reference method is built for field exploration by minimizing the variance of the Kriging model using displacement of the agents to areas of high uncertainty, the proposed criterion (22) only focuses on exploring areas where the maximum could be located. These simulation results support the use of the proposed criterion to limit the exploration area for a faster convergence to the maximum with few information.

6. CONCLUSION

In this paper, a novel method of maximum seeking for a MAS has been presented. It relies on Kriging to perform global optimization of some unknown field. Unlike other methods based on gradient climbing, the method can deal with multiple extrema. The proposed criterion for selecting next sampling points prevents from exploring entirely the search domain, by favoring sampling in areas the more likely to improve the estimate of the global maximum. The dynamics of the agents is also taken into account in the sampling point selection process. Simulation results have been provided for comparison with a reference method, which shows that the number of measurements to be taken by the agents tends to be smaller to get a quick convergence to the maximum of the field. In future work, the potential presence of outliers should be taken into account in the Kriging process to get a more robust maximum seeking scheme.

ACKNOWLEDGEMENTS

This work was partially supported by the DGA, of the French Defence Ministry, by the Paris-Saclay Icode institute, and the Institut Universitaire de France.

REFERENCES

Cheah, C.C., Hou, S.P., and Slotine, J.J. (2009). Region-based shape control for a swarm of robots. *Automatica*, 45(10), 2406–2411.

Choi, J. and Horowitz, R. (2007). Cooperatively learning mobile agents for gradient climbing. In *Proceedings of the 46th IEEE Conference on Decision and Control, New Orleans, LA, USA*, 3139–3144.

Choi, J., Lee, J., and Oh, S. (2008). Swarm intelligence for achieving the global maximum using spatio-temporal gaussian processes. In *Proceedings of the American Control Conference, Seattle, WA, USA*, 135–140.

Choi, J., Oh, S., and Horowitz, R. (2009). Distributed learning and cooperative control for multi-agent systems. *Automatica*, 45(12), 2802–2814.

Cortes, J. (2009). Distributed kriged kalman filter for spatial estimation. *IEEE Transactions on Automatic Control*, 54(12), 2816–2827.

Cox, D.D. and John, S. (1997). SDO: A statistical method for global optimization. *Multidisciplinary Design Optimization: state of the art*, 315–329.

Gu, D. and Hu, H. (2012). Spatial gaussian process regression with mobile sensor networks. *IEEE Transactions on Neural Networks and Learning Systems*, 23(8), 1279–1290.

Jones, D.R., Perttunen, C.D., and Stuckman, B.E. (1993). Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Applications*, 79(1), 157–181.

Jones, D. (2001). A taxonomy of global optimization methods based on response surfaces. *Journal of Global Optimization*, 21(4), 345–383.

Karasev, V. (2012). Controlled Recognition Bounds for Visual Learning and Exploration. *Advances in Neural Information Processing Systems*, 2915–2923.

Kushner, H.J. (1962). A versatile stochastic model of a function of unknown and time varying form. *Journal of Mathematical Analysis and Applications*, 5(1), 150–167.

Ogren, P., Fiorelli, E., and Leonard, N.E. (2004). Cooperative control of mobile sensor networks: Adaptive gradient climbing in a distributed environment. *IEEE Transactions on Automatic Control*, 49(8), 1292–1302.

Schonlau, M. (1997). *Computer Experiments and Global Optimization*. PhD thesis, University of Waterloo, Canada.

Schonlau, M. and Welch, W.J. (1996). Global optimization with nonparametric function fitting. *Proceedings of the ASA, Section on Physical and Engineering Sciences*, 183–186.

Vazquez, E. and Bect, J. (2010). Convergence properties of the expected improvement algorithm with fixed mean and covariance functions. *Journal of Statistical Planning and Inference*, 140(11), 3088–3095.

Wang, P.K.C. (1991). Navigation strategies for multiple autonomous mobile robots moving in formation. *Journal of Robotic Systems*, 8(2), 177–195.

Williams, R.K. and Sukhatme, G.S. (2012). Probabilistic spatial mapping and curve tracking in distributed multi-agent systems. In *Proceedings of the IEEE International Conference on Robotics and Automation, St Paul, MN, USA*, 1125–1130.

Xu, Y., Choi, J., and Oh, S. (2011). Mobile sensor network navigation using gaussian processes with truncated observations. *IEEE Transactions on Robotics*, 27(6), 1118–1131.

Zhang, F. and Leonard, N. (2010). Cooperative filters and control for cooperative exploration. *IEEE Transactions on Automatic Control*, 55(3), 650–663.