

Optimotaxis: A Stochastic Multi-agent Optimization Procedure with Point Measurements^{*}

Alexandre R. Mesquita, João P. Hespanha, and Karl Åström

Center for Control, Dynamical Systems and Computation
University of California, Santa Barbara, CA 93106

Abstract. We consider the problem of seeking the maximum of a scalar signal using a swarm of autonomous vehicles equipped with sensors that can take point measurements of the signal. Vehicles are not able to measure their current position or to communicate with each other. Our approach induces the vehicles to perform a biased random walk inspired by bacterial chemotaxis and controlled by a stochastic hybrid automaton. With such a controller, it is shown that the positions of the vehicles evolve towards a probability density that is a specified function of the spatial profile of the measured signal, granting higher vehicle densities near the signal maxima.

1 Introduction

This paper addresses the problem of controlling a team of autonomous vehicles so as to find the maximum of a scalar function defined over a region of interest, without position and gradient measurements. In the stochastic framework adopted, our goal is to enforce a probability density for the vehicles' positions whose maximum coincides with the maximum of the scalar function. This is achieved by inducing the vehicles into a random motion that mimicks the chemotactic behavior observed in the bacterium *Escherichia coli*. Being unable to directly sense chemical gradients because of its reduced dimensions, this organism is still able to follow the gradient of a chemical attractant, despite the rotational diffusion that constantly changes the bacterium orientation. This is accomplished by switching between two alternate behaviors known as *run* and *tumble* [1, 2].

The problem of finding the maximum of a spatially defined function by moving agents (that we also call vehicles) is often called *source-seeking*. This terminology refers to a specific application of this problem, which consists of finding the source of a chemical substance that is being produced at one particular location, but spreads over a region through a diffusion process. Potential applications for source-seeking include chemical plant safety, hydrothermal vent prospecting and pollution and environmental monitoring.

^{*} This material is based upon work supported by the Inst. for Collaborative Biotechnologies through grant DAAD19-03-D-0004 from the U.S. Army Research Office. The first author was partially funded by CAPES (Brazil) grant BEX 2316/05-6.

We are interested in source-seeking under very limited sensing by the vehicles. Gradient information is often not directly available, either because of noisy and turbulent environments or because the vehicle size is too small to provide accurate gradient measurements, challenges also faced by *E. coli*. In addition, dispensing position measurements is necessary in applications for which inertial navigation systems are expensive, GPS is not available or not sufficiently accurate (as in underwater navigation or cave exploration), or the vehicles are too small or weight-constrained to carry this type of equipment.

Classical techniques from numerical optimization have been adapted for single- and multi-vehicle search strategies when gradients are not explicitly available [3–5]. In [3], the local gradient is estimated by means of a circular movement. The simplex method is implemented with a network of autonomous vehicles in [4]. However, this approach requires the ability to measure the vehicles relative position. Mayhew et al. [5] proposed a hybrid control algorithm to perform a conjugate directions method without position measurements. Control algorithms for networks of vehicles inspired by collective behavior such as fish schooling and chemotaxis are designed in [6, 7]. An extremum seeking strategy is adopted in [8]. Statistical approaches have also been proposed for the case of turbulent environments, but assuming the availability of vehicle position measurements [9, 10]. In general, when convergence is proven in the above mentioned references, this is done exclusively under the assumption that the signal spatial profile is quadratic-like. Bio-inspired techniques have a strong appeal in optimization. Examples are the well-known genetic algorithms and the solutions for the traveling salesman problem inspired by ant colonies [11]. Mimicking chemotaxis is not a new approach to the source-seeking problem, see e.g. [6, 10, 12, 13].

In the *E. coli*, chemotaxis consists of an alternation of two modes of motion called the *run* and *tumble* phases. In the run phase the bacterium swims with constant velocity by rotating its flagella in the counter-clockwise direction. In the tumble phase, by rotating its flagella in the clockwise direction, the bacterium spins around without change in its position and in such a way that it enters the next run phase with arbitrary orientation. Berg and Brown [1] observed that the only parameter that is affected by the concentration of a chemical attractant is the duration of runs. Roughly speaking, the less improvement the bacterium senses in the concentration of the attractant during the run phase, the more probable a direction change (tumble) becomes. Such a motion leads to a distribution whose peak usually coincides with the optimum of the sensed quantity. Whereas the previous works on chemotaxis-based source-seeking rely on a heuristic approach, we present a technique that allows one to control the probability density of the vehicle’s position to permit the estimation of not only the peak of the sensed signal, but also the whole spatial profile of that signal. In fact, with the algorithm proposed, the probability density of the vehicle’s position converges to a pre-specified function of the signal spatial profile. We adopt the suggestive name of *optimotaxis* to designate this search procedure. An important feature of optimotaxis is that it can be used with a broad class of signal profiles, including the ones with multiple maxima, a feature that is shared

with a few other stochastic optimization algorithms which are not constrained by vehicle kinematics [14, 15].

Optimotaxis is motivated by scenarios in which source-seeking is to be solved by a team of autonomous agents. Since the probability density of the agents' position is guaranteed to converge to a pre-specified function of the sensed signal, a supervisor that can sense the positions of the individual agents will be able to monitor the profile of the sensed signal and discern the location of its optimum. Our convergence results (cf. Theorems 1–2) show that by averaging the observations of the position of a single agent over time one recovers the signal profile. In fact, optimotaxis can be performed effectively by a single vehicle. However, given that optimotaxis was specially designed for use with small and cheap vehicles, one should take advantage of the more accurate estimates obtained with multiple vehicles, as we explain in Section 4.

2 Problem Description and Controllers

We consider vehicles with position $\mathbf{x} \in \Omega := \mathbb{R}^d$ and velocity $\mathbf{v} \in V := \rho\mathbb{S}^{d-1}$, where $\rho > 0$, $d \in \{2, 3\}$, and \mathbb{S} is the unity sphere equipped with the Lebesgue measure $d\mu$. We denote by $F(x)$ the scalar function describing the intensity of the sensed signal at position $x \in \Omega$. We define $p(x, v, t)$ as the probability density of finding a vehicle at position x with velocity v at time t . For each fixed time t , $p(x, v, t) \in L^1(\Omega \times V)$, where $\Omega \times V$ is equipped with the product measure $dx \otimes d\mu$.

The objective of optimotaxis is to design a control law to select the velocity $\mathbf{v}(t)$ as a function of the measurements $\{F(\mathbf{x}(\tau)); 0 \leq \tau \leq t\}$ collected up to time t such that the marginal $\int_V p(x, v, t) d\mu(v)$ converges to some density $Q(F(x))$. The function $Q(\cdot)$ is a design parameter called the shaping function which is used to guarantee that $Q(F)$ is a valid probability density function and perhaps to accentuate the maxima of the sensed signal. For example, if $F(x) = 1 - \|x\|^2$, a valid nonnegative shape would be attained with

$$Q(F) = \begin{cases} cF & , \text{ if } F > \delta \\ c\delta e^{F-\delta} & , \text{ if } F \leq \delta \end{cases} , \quad (1)$$

where $\delta > 0$ and c is a normalizing constant. Alternatively, if one is mainly interested in the position of the maxima of $F(x)$, a possible choice would be $Q(F)$ equal to some power of F , which would make maxima more distinct.

Two different controllers inspired by bacterial chemotaxis are presented below.

2.1 Run and Tumble Controller

The jump-diffusion framework that we adopt for optimotaxis was introduced for bacterial chemotaxis in [16, 17]. We consider here a vehicle moving with constant velocity between tumbles. The conditional probability that a tumble does not

occur between the time instants t and s , given that the vehicle was at position x with velocity v at time s is given by

$$\exp\left(-\int_s^t \lambda(x + \tau v, v) d\tau\right), \quad (2)$$

where the positive function $\lambda(\cdot, \cdot) \in L^\infty(\Omega \times V)$ is the tumbling rate. At each tumble, the velocity jumps to a random value $\mathbf{v} \in V$ with probability density $T_{\mathbf{v}^-}$ that may depend on the velocity \mathbf{v}^- just before the tumble. In the bacterial case, the quantities λ and T_v were characterized by Alt [2]. In our work, these should be viewed as control parameters which we will select and that may depend on x and v through the measurements $\{F(\mathbf{x}(\tau)); 0 \leq \tau \leq t\}$.

This controller for optimotaxis is captured by several stochastic hybrid system models that appeared in the literature, including Piecewise-Deterministic Markov Processes (PDPs) [18], our stochastic hybrid models discussed in [19], or the hybrid models initially proposed in [20] by Hu, Lygeros, and co-workers and further expanded in a series of subsequent papers. Fig. 1 depicts a schematic representation of our hybrid model for optimotaxis.




Fig. 1. Hybrid automaton for optimotaxis

The probability density $p(x, v, t)$ was shown to satisfy the following integro-differential equation [16, 17]:

$$\frac{\partial p}{\partial t}(x, v, t) + v \cdot \nabla_x p(x, v, t) = -\lambda p(x, v, t) + \int_V T_{v'}(v) \lambda(x, v') p(x, v', t) d\mu(v') . \quad (3)$$

We note that (3) has been indifferently used in the literature for p as a probability density or simply the density of individuals. When p is regarded as the density of individuals, (3) has a simple intuitive interpretation. On the left-hand side we find a drift term $v \cdot \nabla_x p$ corresponding to the vehicles straight runs, on the right-hand side we find an absorption term $-\lambda(x, v)p(x, v)$, that corresponds to vehicles leaving the state (x, v) , and an integral term corresponding to the vehicles jumping to the state (x, v) . Equation (3) is also known in linear transport theory, where it models the particular case of pure scattering [21, 22]. In that framework, the absorption and the integral terms account for elastic collisions between particles.

2.2 Diffusion Controller

In the limit as the tumbling rate and the velocity approach infinity the jump-diffusion process may approach a pure diffusion process [17]. Inspired by this

observation, we present in this section an alternative controller for optimotaxis that requires the vehicle to turn constantly. This could be applied when the vehicles are capable of relatively high tumbling rates. We do not claim that the model in this section is the diffusion approximation for the model in the run and tumble controller, though it seems to have the same desirable properties of the diffusion approximation. For simplicity we adopt $d = 2$. Let $\mathbf{x} = [\mathbf{x}_1 \ \mathbf{x}_2]^\top \in \Omega$ be the position vector in the plane and θ be such that $\mathbf{v} = [\rho \cos \theta \ \rho \sin \theta]^\top \in V$. We propose a controller given by the following stochastic differential equation:

$$\begin{aligned} d\mathbf{x}_1 &= \rho \cos \theta dt \\ d\mathbf{x}_2 &= \rho \sin \theta dt \\ d\theta &= \sigma(\mathbf{x}, \theta) d\mathbf{w} \ , \end{aligned} \tag{4}$$

where $\mathbf{w}(t)$ is a continuous Wiener process with $E d\mathbf{w}^2 = 1$ and $\sigma(x, \theta)$ is the turning intensity, which we select to attain the desired behavior. The rationale for this controller involves the same original idea of having a turning intensity proportional to the improvement in the measurements. This model with continuous turning allows us to write the Fokker-Planck equation for the probability density of agents $p(x, v, t)$ [23]:

$$\frac{\partial p}{\partial t}(x, v, t) + v \cdot \nabla_x p(x, v, t) = \frac{1}{2} \frac{\partial^2}{\partial \theta^2} (\sigma^2 p(x, v, t)) \ . \tag{5}$$

3 Control law

To obtain the control law we substitute the desired stationary density in equations (3) or (5) for the evolution of the density and then solve for λ and T_v or σ . We will subsequently verify convergence to the desired density. For simplicity, we set the desired stationary density to be independent of v such that $p(x, v, t) = Q(F(x))$. In the following we write $Q(F(x))$ simply as $Q(x)$.

3.1 Run and Tumble Controller

As we substitute $p(x, v, t) = Q(x)$ in the steady state version of (3), we obtain

$$v \cdot \nabla_x Q(x) = -\lambda(x, v) Q(x) + Q(x) \int_V T_{v'}(v) \lambda(x, v') d\mu(v') \ . \tag{6}$$

At this point we need to make the assumption that $Q(x) > 0$ for all $x \in \Omega$. This important assumption will be made throughout this paper. Dividing (6) by $Q(x)$ and rearranging the terms we conclude that we must have

$$\lambda(x, v) = \int_V T_{v'}(v) \lambda(x, v') d\mu(v') - v \cdot \nabla_x \ln Q(x) \ . \tag{7}$$

In the case of a uniformly distributed velocity jump, namely,

$$T_{v'}(v) = \frac{1}{\mu(V)} \ , \tag{8}$$

it is straightforward to solve for $\lambda(x, v)$:

$$\lambda(x, v) = \eta(x) - v \cdot \nabla_x \ln Q(x) \quad , \quad (9)$$

where $\eta(x) = \int_V T_{v'}(v) \lambda(x, v') d\mu(v')$ is some function chosen by the designer to depend on x only through $F(x)$.

This control law is implementable using the past measurements $\{F(\mathbf{x}(\tau)); 0 \leq \tau \leq t\}$ since the tumbling rate depends only on the projection of the gradient in the direction of motion. In fact,

$$\mathbf{v}(t) \cdot \nabla_x Q(\mathbf{x}(t)) = \frac{dQ(\mathbf{x}(t))}{dt} \quad . \quad (10)$$

Notice that $\eta(x)$ is the average tumbling rate at the position x and it must be chosen larger than or equal to $\rho \sup \|\nabla_x \ln Q(x)\|$ to make sure that the tumbling rate λ is positive. In this paper we only consider finite tumbling rates. Thus, here comes a second important restriction on $Q(x)$: $\|\nabla_x \ln Q(x)\|$ must be uniformly bounded. When this happens, we can take η to be independent of x and, according to (2), the probability of an agent maintaining a run with the same direction during the interval $[0, t]$ is given by

$$\begin{aligned} \exp\left(-\int_0^t \lambda(\mathbf{x}(\tau), \mathbf{v}(\tau)) d\tau\right) &= \exp\left(-\int_0^t \eta - \frac{d}{d\tau}(\ln Q(\mathbf{x}(\tau))) d\tau\right) \\ &= e^{-\eta t} \frac{Q(\mathbf{x}(t))}{Q(\mathbf{x}(0))} \quad . \quad (11) \end{aligned}$$

This provides a simple and useful expression for the practical implementation of the search procedure: Suppose that an agent tumbled at time t_k , at that time pick a random variable P uniformly distributed in the interval $[0, 1]$ and tumble when the following condition holds

$$Q(\mathbf{x}(t)) \leq P e^{\eta(t-t_k)} Q(\mathbf{x}(t_k)), \quad t \geq t_k \quad . \quad (12)$$

As opposed to what (10) seems to imply, one does not need to take derivatives to implement (9). Also, the control law is not changed if a constant scaling factor is applied to $Q(x)$. It is important to remark that η may be adjusted online. An agent may begin a search with $\eta = \epsilon > 0$ and if at some time t it observes that $\eta < \bar{\eta} = t^{-1} \ln Q(\mathbf{x}(t))/Q(\mathbf{x}(0))$, then it updates η to $\bar{\eta} + \epsilon$. The use of a small residue ϵ grants a positive λ . In this case, one can prove that the probability to have the vehicle visiting any neighborhood in space is positive. Hence, η will eventually converge to $\rho \sup \|\nabla_x \ln Q(x)\| + \epsilon$. A more elaborate adaptation can be obtained by having $\eta(x)$ as a function of the measurements $F(x)$, which would reduce the number of unnecessary tumbles.

We note that most physical quantities propagate with spatial decay not faster than exponential, which allows for the uniform boundedness of $\|\nabla \ln Q(x)\|$. If, however, $F(x)$ has a faster decay rate, it may still be possible to achieve boundedness of $\|\nabla_x \ln Q(x)\|$ via reshaping (i.e. selecting an appropriate Q) as long as $F(x) > 0$ for all $x \in \Omega$ and one knows its maximum decay rate.

3.2 Diffusion Controller

Substituting $p(x, v, t) = Q(x)$ in (5) and integrating twice in θ one obtains:

$$v \cdot \nabla_x Q(x) + \frac{1}{2} \sigma^2(x, \theta) Q(x) = \theta c_1(x) + c_2(x) . \quad (13)$$

Solving for $\sigma^2(x, \theta)$ we have

$$\sigma^2(x, \theta) = \theta d_1(x) + d_2(x) - 2v \cdot \nabla_x \ln Q(x) . \quad (14)$$

We set $d_1(x) = 0$ since there is no advantage in having $\sigma^2(x, \theta)$ growing linearly with θ . Thus, we can rewrite (14) as

$$\sigma^2(x, v) = \eta(x) - 2v \cdot \nabla_x \ln Q(x) . \quad (15)$$

Again, the only information needed to implement $\sigma^2(x, v)$ is the measurements $F(\mathbf{x}(t))$ collected along the vehicle's trajectory. As before, boundedness of $\|\nabla_x \ln Q(x)\|$ is an important condition and $\eta(x)$ may be adjusted online. A simple implementation of (4) with $\sigma(x, v)$ given by (15) can be done using Euler's approximation.

The addition of exogenous white Gaussian noise to $\dot{\theta}$ in (4) leads to the addition of a positive constant to $\sigma^2(x, v)$ in (5), which is equivalent to an increased η . Thus, this kind of disturbance does not change the stationary density. The same conclusion is true if we consider white Gaussian noise added to $\dot{\rho}$ in (5). Therefore, an important property of this controller is that the stationary density is robust to additive white Gaussian noise applied in the vehicle's bodyframe.

4 Convergence to the steady-state

In this section we analyze the convergence of solutions to the desired stationary density $Q(x)$. We consider mild solutions [24] to the initial value problem defined by (3) and an initial density $p(x, v, 0) \in \mathcal{D} := \{f \in L^1(\Omega \times V); f \geq 0, \|f\| = 1\}$. The main result is stated in Theorem 1, which implies that $Q(x)$ can be estimated through the time average of the observed vehicles' position. In particular, with a proper choice of $Q(x)$, the maximum of $F(x)$ will be located in the neighborhood that is most often visited by the vehicle.

Theorem 1. *Assume that $Q(x) > 0$ and $\|\nabla_x \ln Q(x)\| \in L^\infty(\Omega)$. If T_v and λ are chosen according to (8) and (9), respectively, such that λ is uniformly bounded and strictly positive, then*

1. *A mild solution $p(x, v, t)$ to (3) exists and is unique for all $t \geq 0$ and all initial densities $p(x, v, 0) \in \mathcal{D}$.*
2. *$Q(x)$ is the unique stationary density for (3).*
3. *For any initial density $p(x, v, 0) \in \mathcal{D}$,*

$$\lim_{N \rightarrow \infty} \left\| \frac{1}{N} \sum_{k=0}^{N-1} p(x, v, k\tau) - Q(x) \right\| = 0 \quad (16)$$

for all $\tau > 0$.

We note that the continuous time average

$$\frac{1}{t} \int_0^t p(x, v, t) dt \quad (17)$$

also converges in norm to $Q(x)$ as $t \rightarrow \infty$, see e.g. [25, Cor. VIII.7.4].

Theorem 1 provides the basis for a procedure to estimate $Q(x)$ by observing the position of N agents: We start by partitioning the region of interest into a family of sets $\{A_i \subset \Omega\}$ and then we sample the vehicles' positions at times $k \in \{0, \tau, 2\tau, \dots, (K-1)\tau\}$ and count the number of times that a vehicle is observed in each set A_i . It turns out that the fraction of samples corresponding to vehicles observed over the set A_i provides an asymptotically correct estimate of the average value of $Q(x)$ on the set A_i . To see why this is the case, we define

$$f_{N,K} = \frac{1}{NK} \sum_{n=0}^{N-1} \sum_{k=0}^{K-1} f(\mathbf{x}_n(k\tau)) , \quad (18)$$

for some $f \in L^\infty(\Omega)$ and $\tau > 0$. Assuming that the agents have mutually independent motion, by the law of large numbers we have that $f_{N,K}$ converges almost surely as $N \rightarrow \infty$. Moreover, by the specific version of the law of large numbers in [26], we also have that $f_{N,K}$ converges almost surely as $K \rightarrow \infty$. In particular, if f is the characteristic function of a measurable set A_i , then

$$f_{N,K} \rightarrow \int_{A_i} Q(x) dx \text{ a.s.} \quad (19)$$

as $K \rightarrow \infty$. This shows that $Q(x)$ can be estimated by averaging the observations of the agents position as in (18). The use of multiple agents ($N > 1$) improves the estimates according to the relation

$$\text{var}(f_{N,K}) = \frac{\text{var}(f_{1,K})}{N} . \quad (20)$$

Proof of Theorem 1 (outline). A complete proof is given in [27]. The existence and uniqueness results rely on previous results on linear transport theory [21] using semigroups of linear operators. To prove the uniqueness of the stationary density we take $q(x, v, t)$ and $r(x, v, t)$ to be two convex combinations of a stationary density $p(x, v, t)$ and $Q(x)$. We consider the Kullback-Leibler divergence between $q(x, v, t)$ and $r(x, v, t)$

$$H(t) = \int_{\Omega} \int_V q(x, v, t) \ln \frac{q(x, v, t)}{r(x, v, t)} dx d\mu(v) , \quad (21)$$

and show that $\dot{H}(t) = 0$ for all $t > 0$ implies $p = Q$. The convergence of the Cesàro averages in part 3 of the theorem is a direct consequence of the mean ergodic theorem provided that $Q(x) > 0$ is the unique invariant density. See e.g. Theorem 5.2.2 of [28] or [29, Chap. 2]. \square

It is worthwhile to remark that $p(x, v, t)$ actually converges to $Q(x)$ in norm if λ is bounded away from zero. The proof of this result is based on a result from [30] that states that, if $V(t)$ is a Harris operator for every $t > 0$ and $Q(x) > 0$ is its unique invariant density, then $V(t)f$ converges to $Q(x)$ in norm for all $f \in \mathcal{D}$. It turns out that one can prove that the semigroup generated by (3) consists of Harris operators [27]. In addition, we conjecture that convergence is exponential. The proof of this result would be possible with a generalization of some results in the spectral theory of linear transport operators in unbounded domains presented in [31]. However, such a generalization is not yet available.

Next, we state the corresponding result for the diffusion controller.

Theorem 2. *Assume that $Q(x) > 0$ and $\|\nabla_x \ln Q(x)\| \in L^\infty(\Omega)$. If σ^2 is chosen according to (15) such that σ^2 is uniformly bounded and strictly positive, then*

1. *A mild solution $p(x, v, t)$ to (5) exists and is unique for all $t \geq 0$ and all initial densities $p(x, v, 0) \in \mathcal{D}$.*
2. *$Q(x)$ is the unique stationary density for (5).*
3. *For any initial density $p(x, v, 0) \in \mathcal{D}$,*

$$\lim_{N \rightarrow \infty} \left\| \frac{1}{N} \sum_{k=0}^{N-1} p(x, v, k\tau) - Q(x) \right\| = 0 \quad (22)$$

for all $\tau > 0$.

The proof of this result is also presented in [27].

5 Numerical Results and Discussion

In this section we present numerical experiments to illustrate the proposed optimization procedure. We adopt preferentially the run and tumble controller with a constant function $\eta(x)$. The results for the diffusion case are similar but with slightly faster convergence. The desired stationary density is taken to be $Q(F(x)) = cF^n(x)$, where c is a normalizing constant and n is an integer.

The main capability of optimotaxis, the localization of the global maximum, is stressed in Fig. 2. We observe a swarm of agents that starts from the upper left corner (1), initially clusters around a local maximum (2) and then progressively migrates to the global maximum (3,4). When the equilibrium is reached, most agents concentrate in a neighborhood of the global maximum. Yet, a portion of the agents clearly indicates the existence of the local maximum. We notice that the center of mass of the swarm goes straight through the local maximum to the global one. This feature is not shared with most deterministic optimization procedures and even with some stochastic ones. As a bonus, the information on secondary sources (local maxima) is not lost.

We use the Kullback-Leibler divergence $H(t)$ (defined in Section 4) between $p(x, v, t)$ and the convex combination $10Q(x)/11 + p(x, v, t)/11$ to analyze the speed of convergence to the desired stationary density. One useful property of




Fig. 2. Different stages of optimotaxis in the presence of two maxima. Black dots represent agents position whereas the background intensity represents the signal intensity. $F(x) = 0.4e^{-\|x\|} + 0.6e^{-\|x-[1.5 \ -1.5]'\|}$, $Q(x) = F^n(x)$ with $n = 10$, and $\rho = 1$.

this Kullback-Leibler divergence is that $H = 0$ iff $p = Q$. Moreover, $\alpha(\|p - Q\|) \leq H(t) \leq \beta(\|p - Q\|)$, where α, β are class \mathcal{K} functions [32]. Using a space grid with resolution 0.068, we calculate $H(t)$, which is shown to converge to zero in Fig. 3.

Also included in Fig. 3 is the evolution of $H(t)$ when the measurements are quantized and when exogenous noise is added. In the quantization case, we used the quantized version of the desired density $Q(x)$ to calculate $H(t)$. Interestingly, the addition of noise does not seem to affect considerably the transient response. Nevertheless, the residual error is greater due to the fact that the stationary density is not the one expected. On the other hand, quantization has a much more negative impact on optimotaxis performance. Yet, we believe that convergence to a quantized $Q(x)$ does happen but at a low speed.

The sensitivity of the procedure with respect to the parameter n of the shaping function is studied with Fig. 4. The mean-square error of the vehicles position with respect to the maximum is used as a performance index. One notices that the performance degrades for n too low or too high. In particular, the sensitivity to noise and quantization increases with n . This suggests that an interesting strategy to reduce the effect of uncertainties and quantization is to assign agents with different values of n . In this case, the observed density would converge to an arithmetic average of the powers $F^n(x)$. Thus, the mean-square error would be smaller than the error corresponding to the maximum or minimum value of the chosen n .




Fig. 3. Evolution of the Kullback-Leibler divergence for: the noiseless case (solid), the quantized measurements case (dashed), and the exogenous noise case (dash-dotted). The number of quantization levels is 128. The noise added to \dot{v} is white Gaussian with standard deviation 10^{-2} along each axis. 100 agents were uniformly deployed in the rectangle $[-2.5, 2.5] \times [-2.5, 2.5] \times V$. Refer to Fig. 2 for more details.

5.1 Chemotaxis and Optimotaxis

It is remarkable that the expression for λ in (9) has the same structure of some biochemical models for the tumbling rate of the *E. coli*; see, for instance, Alt [2, Equation 4.8]. This author essentially proposed the existence of a chemical activator for the locomotion mechanism such that a tumble would occur each time the concentration of this activator would become less than a certain value. The concentration of this activator would jump to a high value at tumbles and decrease at a rate corresponding to η in (9). A receptor-sensor-mechanism would then regulate the additional generation of the activator (this corresponds to $v \cdot \nabla_x Q(x)$ in (9)), which would modulate the run length. We find surprising that our reverse engineering design resulted in an expression for λ similar to the one in bacterial chemotaxis. In fact, though the use of tumble and run in optimotaxis is inspired by chemotaxis, one would not necessarily expect that our choice of the tumbling rate would lead to control laws similar to the biochemical models in bacteria. This fact suggests that the bacteria evolutionary process might have selected a taxis mechanism with the same desirable properties of our model.

To understand what these desired properties might be, let us suppose that bacteria are performing optimotaxis as it is described in this paper. Let $p(x, v, t)$ be the spatial density of bacteria and let $Q(x)$ be some function related to the concentration of nutrients. From Section 4, we know that $p(x, v, t)$ converges in norm, which implies that

$$H(t) = - \int_{\Omega} \int_V p(x, v, t) \ln \left(\frac{1}{2} + \frac{1}{2} \frac{Q(x)}{p(x, v, t)} \right) dx d\mu(v) \rightarrow 0 . \quad (23)$$




Fig. 4. Mean-square error with respect to the maximum of $F(x) = e^{-\|x\|}$ as a function of n . Noiseless case (solid), quantized $F(x)$ (dashed), and exogenous noise (dash-dotted). The number of quantization levels is 128. The noise added to \dot{v} is white Gaussian with standard deviation 10^{-3} in each axis. $\rho = 1$.

Thus, $H(t)$ can be regarded as a cost functional that is being minimized by optimotaxis/chemotaxis. More specifically, we notice that what is being maximized is the expected value of a concave function of Q/p , which is the ratio of the concentration of nutrients per the density of organisms. This is a meaningful cost for the population of bacteria as a whole.

It is important to remark that the jump velocity probability density function T_v is not uniform in bacteria. This supports our belief that better results for optimotaxis might be obtained with more sophisticated choices for T_v .

6 Conclusion

A random optimization algorithm based on bacterial chemotaxis was presented. This algorithm is mainly intended for application in a swarm of agents whose mission is to find the maximum of a measured quantity. The most attractive features of the procedure are its simplicity and low cost of implementation as well as the identification of both global and local maxima. The only measurement needed by the agents is the signal of interest and the only information that must be known a priori is a bound on the spatial decay of the measured quantity. The convergence of the agents probability density to a specified function was demonstrated. Some robustness to exogenous disturbance was also demonstrated. In addition, insight was gained on what bacterial chemotaxis might try to optimize.

We note that the proposed choice of control parameters may not be optimal. Hence, alternative choices for the tumbling rate and for the probability density of velocity jumps are important themes for further investigation. Future research directions also include the study of optimotaxis in compact spatial domains

and signal spatial profiles with discontinuities. Alternatively, it may be worth to explore the application of these ideas to solve problems in numerical optimization by arrays of independent processors.

Acknowledgments

The authors are thankful to Professors Mustafa Khammash, Upamanyu Madhok, and Prashant Mehta for valuable comments.

References

1. Berg, H., Brown, D.: Chemotaxis in *Escherichia coli* analysed by three-dimensional tracking. *Nature* **239**(5374) (October 1972) 500–504
2. Alt, W.: Biased random walk models for chemotaxis and related diffusion approximations. *J Math Biol* **9**(2) (April 1980) 147–177
3. Burian, E., Yoerger, D., Bradley, A., Singh, H.: Gradient search with autonomous underwater vehicles using scalar measurements. In: Proceedings of the 1996 Symposium on Autonomous Underwater Vehicle Technology, 1996. AUV '96. (June 1996) 86–98
4. Sousa, J., Johansson, K., Silva, J., Speranzon, A.: A verified hierarchical control architecture for co-ordinated multi-vehicle operations. *International Journal of Adaptive Control and Signal Processing* **21**(2-3) (September 2007) 159–188
5. Mayhew, C., Sanfelice, R., Teel, A.: Robust source-seeking hybrid controllers for autonomous vehicles. In: Proceedings of the 2007 American Control Conference. ACC '07. (July 2007) 1185–1190
6. Hoskins, D.A.: Least action approach to collective behavior. In Parker, L.E., ed.: Proc. SPIE Vol. 2593, p. 108–120, *Microrobotics and Micromechanical Systems*, Lynne E. Parker; Ed. Volume 2593 of Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference. (December 1995) 108–120
7. Bachmayer, R., Leonard, N.: Vehicle networks for gradient descent in a sampled environment. In: Proceedings of the 41st IEEE Conference on Decision and Control. (December 2002) 112–117
8. Zhang, C., Arnold, D., Ghods, N., Siranosian, A., Krstic, M.: Source seeking with non-holonomic unicycle without position measurement and with tuning of forward velocity. *Systems & Control Letters* **56**(3) (March 2007) 245–252
9. Pang, S., Farrell, J.: Chemical plume source localization. *IEEE Transactions on Systems, Man and Cybernetics, Part B* **36**(5) (October 2006) 1068–1080
10. Vergassola, M., Villermaux, E., Shraiman, B.: 'infotaxis' as a strategy for searching without gradients. *Nature* **445**(7126) (January 2007) 406–409
11. Dorigo, M., Gambardella, L.M.: Ant colony system: A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation* **1**(1) (April 1997) 53–66
12. Ferrée, T., Lockery, S.: Computational rules for chemotaxis in the nematode *C. Elegans*. *Journal of Computational Neuroscience* **6** (1999) 263–277
13. Dhariwal, A., Sukhatme, G.S., Requicha, A.A.: Bacterium-inspired robots for environmental monitoring. In: IEEE International Conference on Robotics and Automation, New Orleans, Louisiana, IEEE (Apr 2004) 1436–1443

14. Aluffipentini, F., Parisi, V., Zirilli, F.: Global optimization and stochastic differential-equations. *Journal of Optimization Theory and Applications* **47**(1) (1985) 1–16
15. Dekkers, A., Aarts, E.: Global optimization and simulated annealing. *Math. Program.* **50**(3) (1991) 367–393
16. Stroock, D.: Some stochastic processes which arise from a model of the motion of a bacterium. *Probability Theory and Related Fields* **28**(4) (December 1974) 305–315
17. Papanicolaou, G.: Asymptotic analysis of transport processes. *Bulletin of the American Mathematical Society* **81**(2) (March 1975) 330–392
18. Davis, M.H.A.: Markov models and optimization. Monographs on statistics and applied probability. Chapman & Hall, London, UK (1993)
19. Hespanha, J.P.: Modeling and analysis of stochastic hybrid systems. *IEE Proc — Control Theory & Applications, Special Issue on Hybrid Systems* **153**(5) (2007) 520–535
20. Hu, J., Lygeros, J., Sastry, S.: Towards a theory of stochastic hybrid systems. In Lynch, N.A., Krogh, B.H., eds.: *Hybrid Systems: Computation and Control*. Volume 1790., Springer (2000) 160–173
21. Kaper, H.G., Lekkerkerker, C.G., Hejtmanek, J.: *Spectral Methods in Linear Transport Theory*. Birkhauser Verlag (1982)
22. Mokhtar-Kharroubi, M.: *Mathematical Topics in Neutron Transport Theory*. World Scientific, Singapore (1997)
23. Risken, H.: *The Fokker-Planck Equation - Methods of Solution and Applications*. Springer Series in Synergetics. Springer-Verlag, Berlin (1984)
24. Pazy, A.: *Semigroup of Linear Operators and Applications to Partial Differential Equations*. Springer-Verlag, New York (1983)
25. Dunford, N., Schwartz, J.: *Linear Operators*. Volume VII of Pure and Applied Mathematics. Interscience Publishers, New York (1957)
26. Derriennic, Y., Lin, M.: Uniform ergodic convergence and averaging along markov chain trajectories. *Journal of Theoretical Probability* **7**(3) (1994) 483–497
27. Mesquita, A., Hespanha, J., Åstrom, K.: Optimotaxis: A stochastic multi-agent optimization procedure with point measurements: Extended version. Technical report, University of California, Santa Barbara (October 2007) Available at <http://www.ece.ucsb.edu/~hespanha/techreps.html>.
28. Lasota, A., Mackey, M.: *Chaos, Fractals, and Noise. Stochastic Aspects of Dynamics*. Volume 97 of Applied Mathematical Sciences. Springer-Verlag, New York (1994)
29. Krengel, U.: *Ergodic Theorems*. Volume 6 of de Gruyter Studies in Mathematics. Walter de Gruyter, Berlin; New York (1985)
30. Pichór, K., Rudnicki, R.: Continuous markov semigroups and stability of transport equations. *Journal of Mathematical Analysis and Applications* **249** (2000) 668–685
31. Mokhtar-Kharroubi, M., Sbihi, M.: Spectral mapping theorems for neutron transport, l^1 -theory. *Semigroup Forum* **72** (2006) 249–282
32. Lin, J.: Divergence measures based on shannon entropy. *IEEE Transactions on Information Theory* **37**(1) (January 1991) 145–151