2013 European Control Conference (ECC)
July 17-19, 2013, Zürich, Switzerland.

Stochastic localization of sources with convergence guarantees

Stephan M. Huck and John Lygeros

Abstract— **We establish convergence guarantees for a recently proposed Markov Chain Monte Carlo (MCMC) method to locate source(s) of a certain concentration field. Our method utilizes a Markovian controller to control the motion of autonomous vehicles on a compact search domain. The distribution of the resulting discrete-time Markov chain is used to estimate the locations of the sources. To guarantee the correctness of the localization, we prove that the existing invariant measure for the Markov chain is unique. The chain is shown to be uniform ergodic and will converge to its stationary distribution. The theoretically derived convergence rate is compared to results from numerical simulations.**

I. INTRODUCTION

The interest in stochastic localization of sources in our previously published work [7] was based on an underwater search scenario, where some fluid, e.g., fresh water, is emitted from a source and gives rise to a measurable concentration field. Instead of applying modelling techniques to capture this complex process and deduce the position of the sources, autonomous vehicles were utilized to explore the region of interest relying only on local measurements. In general the aforementioned problem is of interest beyond marine applications, since the framework can easily be related to applications with different scenarios and other types of agents. For example, one can think of search and rescue missions for drones or mobile robots with the similar goal of locating the maximum of a certain beacon signal. In our previous work we presented a MCMC like algorithm constructing an estimate of the concentration signal, taking into account unicycle dynamics and non-stopping constraints for the vehicles.

There exist several approaches that address the described problem using autonomous vehicles. Some make use of deterministic gradient methods, for example, [1], [3], [8], [15]. The methods presented in [11], [12] utilize multiple agents to drive the group in a formation towards the maxima. A second class of approaches to this problem is relying on stochastic techniques. In [2] simulated annealing is carried out with a single vehicle. The authors of [5] build a likelihood estimate of the source location based on a hidden Markov method. For a more extensive literature review the reader is referred to [7].

The stochastic method described in [9], which motivated our initial work, uses the vehicles to randomly explore the space. Observing the probability distribution of the vehicles gives rise to an estimate of the unknown concentration field. Similarly, our method relies on the analysis of the marginal probability distribution of the underlying Markov chain. Therefore, an important aspect to guarantee that the applied MCMC algorithm leads to reliable results, is whether the Markov chain has a *unique* stationary distribution and if the process will *converge* to it. In this work, we will extend the theoretical framework of our method by addressing these questions. First we introduce a more realistic setting in which the dynamics are effected by noise and show that the previously established existence of an invariant measure still holds. Second we prove uniqueness of the existing stationary distribution and give a convergence guarantee for our method by proving that the Markov chain is uniformly ergodic. Furthermore we are able to obtain a theoretical bound on the convergence rate.

The remainder of this paper is organized as follows. A brief description of the problem setup is given in Section II. In Section III we restate the Markov Chain formulation together with the algorithm proposed in [7]. We provide the existence, uniqueness and ergodicity results in Sections IV, V and VI respectively. The convergence results from simulation are given in Section VII and a discussion about future directions can be found in Section VIII.

II. PROBLEM SETUP

A. Dynamics under noise

We consider vehicles moving in a plane $\mathcal{X} \subset \mathbb{R}^2$ with discrete time unicycle dynamics. The position $x \in \mathcal{X}$ is affected by noise and the control input u to the system is applied via the heading angle $\theta \in [-\pi, \pi]$ according to:

$$\begin{pmatrix} x_{k+1} \\ \theta_{k+1} \end{pmatrix} = \begin{pmatrix} x_k + v(\theta_k)T \\ \theta_k \end{pmatrix} + \begin{pmatrix} w \\ u_k \end{pmatrix}. \tag{1}$$

The noise is given by $w \sim \mathcal{U}\{w \in \mathbb{R}^2 |\ ||w|| \leq r_D\}$ with bound $r_D < \bar{v}T$ and $v(\theta) = \bar{v} \begin{pmatrix} \cos(\theta) \\ \sin(\theta) \end{pmatrix}$ for a constant speed $\bar{v}$. Note that we assume the sampling time T to be big enough for a low-level navigation function to realize every change of orientation. Here $\sim \mathcal{U}$ means, that a random variable is uniformly distributed over a set. Furthermore, each agent is capable of measuring a certain normalized concentration function $C : \mathcal{X} \to [0, 1]$ at its current location.

The initial position is extracted on a subset $\tilde{\mathcal{X}}$ of the domain $\mathcal{X}$ as $x_0 \sim \mathcal{U}\{x \in \tilde{\mathcal{X}} | \tilde{\mathcal{X}} \subseteq \mathcal{X}\}$ where the subsequent extraction of $\theta_0 \sim \mathcal{U}\{\theta \in [-\pi, \pi] | x_0 + v(\theta)T \in \mathcal{X}_r\}$ guarantees that the vehicle does not leave $\mathcal{X}$ after the first step. Note that in contrast to [7], the orientation is now

The authors are with the Automatic Control Laboratory, Department of Information Technology and Electrical Engineering, ETH Zurich, Zurich 8092, Switzerland. Email: {huck,lygeros}@control.ee.ethz.ch

This research was partially supported by the NanoTera-SSSCT grant under the project NetCam

constrained by the set $\mathcal{X}_r \subset \mathcal{X}$ whose boundary $\partial\mathcal{X}_r$ has the distance of r_D to $\partial\mathcal{X}$ everywhere.

III. MARKOV CHAIN FORMULATION

We briefly restate here the controller setup, the resulting Markov chain and algorithm for the dynamics in (1). For a detailed discussion on the basic setup, the reader is referred to [7].

Consider the measurable space $(\mathcal{S}, \mathcal{B}(\mathcal{S}))$ with $\mathcal{B}(\mathcal{S})$ being the Borel σ-algebra on the metric space $\mathcal{S} \coloneqq \bigcup_{x\in\mathcal{X}} x \times \vartheta(x)$. The set of angles $\vartheta(x)$ for every $x \in \mathcal{X}$ is given by

$$\vartheta(x) \subseteq [-\pi, \pi] \quad \text{with} \quad \vartheta(x) = \{\theta | x + v(\theta)T \in \mathcal{X}_r\}. \quad (2)$$

We use a Markovian controller with an accept-reject mechanism to compute the input in each step as

$$u_k = \begin{cases} \tilde{\theta}_{k+1} \sim \mathcal{U}\{\theta \in \Theta(s_k)\} & \text{w. p.} \quad \alpha(s_k) \\ 0 & \text{w. p.} \quad 1 - \alpha(s_k), \end{cases} \quad (3)$$

where

$$\Theta(s_k) \triangleq \{\theta \in [-\pi, \pi] \mid x_k + v(\theta_k) + v(\theta_k + \theta)T \in \mathcal{X}_n\}. \quad (4)$$

In essence, a change of orientation is proposed in each step and if it is rejected, the vehicle keeps moving in the previous direction.

Similar to $\mathcal{X}_r$ above, the boundary $\partial\mathcal{X}_n$ has the distance of $2r_D$ to $\partial\mathcal{X}$ everywhere. When expressing the set $\mathcal{X}$ through the Minkowski sum $\mathcal{X}_n \oplus B = \mathcal{X}$, where B is a ball of radius $2r_D$, it is clear that discontinuities of the border of $\mathcal{X}$ are not allowed. Hence, we require the following assumption to ensure existence of $\mathcal{X}_n$ and $\mathcal{X}r$.

Assumption 1: $\mathcal{X}$ is a compact convex set with non-empty interior and has a boundary which is differentiable everywhere. The curvature[1] of the boundary is always smaller than $\frac{1}{2r_D}$.

In (3) the *acceptance criterion* $\alpha : \mathcal{S} \to [0, 1]$ can be any function with the following functionalities, which are justified in detail in [7]:

1) α is monotonically increasing in C.
2) α depends continuously on s.
3) $\alpha = 1$ near the boundary of the set $\mathcal{X}$, possibly also angle dependent.

The closed-loop motion of the system (1) with control (3) generates a discrete-time Markov chain $\{s_k\}_{k\geq 0}$, where $s_k \coloneqq (x_k, \theta_k) \in \mathcal{S}$. An outline of the resulting process is given in Algorithm 1. The transition kernel of the Markov chain $\{s_k\}_{k\geq 0}$ becomes

$$\begin{aligned} \mathcal{P}(s_k, ds) &= p(s_k, s)ds \\ &= \left[\frac{\mathbb{I}_{\{x\in D(s_k)\}}}{\gamma_{\mathcal{X}}(s_k)} \frac{\mathbb{I}_{\{\theta\in\Theta(s_k)\}}}{\gamma_{\Theta}(s_k)} \alpha(s_k) \right. \\ &\quad \left. + \frac{\mathbb{I}_{\{x\in D(s_k)\}}}{\gamma_{\mathcal{X}}(s_k)} \delta_{\theta_k}(\theta)(1 - \alpha(s_k)) \right] ds, \end{aligned} \quad (5)$$

[1]For the definition of the curvature see [4, Chapter 1-5]

Algorithm 1 MCMC AUV Algorithm

Require: Initial state $s_0 = (x_0, \theta_0)$

set $k = 0$
loop
  update $x_{k+1} = x_k + v(\theta_k)T + w$
  generate proposal angle $\tilde{\theta}_{k+1}$
  calculate the acceptance probability $\alpha(s_k)$
  update

$$\theta_{k+1} = u_k = \begin{cases} \tilde{\theta}_{k+1} & \text{w. p.} \quad \alpha(s_k) \\ \theta_k & \text{w. p.} \quad 1 - \alpha(s_k) \end{cases}$$

  set $k = k + 1$
end loop

where $\mathbb{I}$ is the usual indicator function, $D(s_k) \subset \mathcal{X}$ is a disc with radius $r_D < \bar{v}T$ centered at $x_k + v(\theta_k)T$ and the normalizing factors are given by

$$\gamma_{\Theta}(s_k) = \int \mathbb{I}_{\{\theta\in\Theta(s_k)\}} d\theta \quad (6)$$

$$\gamma_{\mathcal{X}}(s_k) = \int \mathbb{I}_{\{x\in D(s_k)\}} dx = \pi r_D^2. \quad (7)$$

Note that by construction of the proposal set for $\tilde{\theta}_{k+1}$ in (4), the factor $\gamma_{\mathcal{X}}$ is constant for all $s_k \in \mathcal{S}$.

IV. EXISTENCE

The existence proof for invariant measures of the Markov chain $\{s_k\}_{k\geq 0}$ follows mainly the same lines as in [7]. However, slight modifications are required since the transition kernel (5) takes the noise affected dynamics into account.

Assumption 2: The curvature of the boundary is never equal to $\frac{1}{\bar{v}T}$. The function α is continuous with respect to s. Both assumptions are technical and can be realized by construction.

Proposition 3: Under Assumptions 1 and 2 the Markov chain $\{s_k\}_{k\geq 0}$ generated by Algorithm 1 with the transition kernel (5) has an invariant measure.

Proof: We use [6, Theorem 7.2.3], i.e., for an invariant measure to exist, it is sufficient to prove the *weak Feller* property for the transition kernel P. That is, we need to show that for a generic continuous and bounded function f on $\mathcal{S}$, the mapping Pf, defined as $(Pf)(\cdot) \coloneqq \int_{\mathcal{S}} P(\cdot, dy)f(y)$, is continuous and bounded. From (5), we have that

$$\begin{aligned} (Pf)(s_k) &= \int_{\mathcal{S}} p(s_k, s) f(s) ds \\ &= \int_{\mathcal{S}} \frac{\mathbb{I}_{\{x\in D(s_k)\}}}{\gamma_{\mathcal{X}}(s_k)} \left[\frac{\mathbb{I}_{\{\theta\in\Theta(s_k)\}}}{\gamma_{\Theta}(s_k)} \alpha(s_k) \right. \\ &\quad \left. + \delta_{\theta_k}(\theta)(1 - \alpha(s_k)) \right] f(x, \theta) dx d\theta \\ &= \int_{\mathcal{S}_D} \frac{\alpha(s_k)}{\gamma_{\mathcal{X}}(s_k)} \frac{\mathbb{I}_{\{\theta\in\Theta(s_k)\}}}{\gamma_{\Theta}(s_k)} f(x, \theta) \end{aligned}$$

$$
\begin{aligned}
&+ \frac{1-\alpha(s_k)}{\gamma_{\mathcal{X}}(s_k)} \delta_{\theta_k}(\theta) f(x,\theta) dx d\theta \\
&= \int_{D(s_k)} \int_{\Theta(s_k)} \underbrace{\frac{\alpha(s_k)}{\gamma_{\mathcal{X}}(s_k)\gamma_{\Theta}(s_k)}}_{c_a(s_k)} f(x,\theta) dx d\theta \\
&+ \int_{D(s_k)} \underbrace{\frac{1-\alpha(s_k)}{\gamma_{\mathcal{X}}(s_k)}}_{c_r(s_k)} f(x,\theta_k) dx, \qquad (8)
\end{aligned}
$$

where $\mathcal{S}_D := \bigcup_{x \in D(s_k)} x \times \vartheta(x)$ and the fact that $\Theta(s_k) \subseteq \vartheta(x)\ \forall x \in D(s_k)$ was used to separate the integration for the first term in (8). The same separation is possible for the second integral, since $\theta_k \in \vartheta(x)$ whenever $\alpha(s_k) \neq 1$. By the curvature part of Assumption 2, $\gamma_{\Theta}(s_k)$ is continuous as the size of the set $\Theta(s_k)$ in (6) changes continuously with respect to s_k. Further, $\gamma_{\mathcal{X}}(s_k)$ is constant by construction (7) and it holds that $\gamma_{\mathcal{X}}(s_k) \neq 0$ and $\gamma_{\Theta}(s_k) \neq 0$. Hence, using Assumption 1 and the functionalities 1)-3) for α in III, the constants c_a and c_r are continuous and bounded with respect to s_k. It is straight forward to show that both integrations in Equation (8) over the sets $D(s_k)$ and $\Theta(s_k)$ are continuous and bounded with respect to s_k [14]. Hence, we conclude that Pf is continuous and bounded and the weak Feller property of the transition kernel P follows.

■

V. UNIQUENESS

First, we make another assumption about the acceptance criterion which is fair, since α is a design choice.

Assumption 4: The acceptance probability α is bounded away from zero, i.e. $\alpha \in [\epsilon, 1]$ for some $\epsilon > 0$.

Given that Proposition 3 holds, we can state that the existing invariant probability measure is unique.

Proposition 5: Under Assumptions 1, 2 and 4 the Markov chain $\{s_k\}_{k \geq 0}$ generated by Algorithm 1 with the transition kernel (5) has an unique invariant measure.

Proof: In view of [6, Proposition 4.2.2] uniqueness of an existing invariant probability measure is achieved by showing *φ-irreducibility*. This is given, if there exists a (non-trivial) measure φ on $\mathcal{B}(\mathcal{S})$ such that, whenever $\varphi(A) > 0$, $A \in \mathcal{B}(\mathcal{S})$, we have $L(s, A) > 0$ for all $s \in \mathcal{S}$ (c.f. definition of φ-irreducibility in [10]). Here $L(s, A) = P(\tau_A < \infty)$ with τ_A being the first return time to set A. By [10, Proposition 4.2.1] this is equivalent to the existence of an $m > 0$ (possibly depending on both s and A) such that $P^m(s, A) > 0$.

In the following we will make use of geometric arguments to achieve φ-irreducibility. By studying the growth of the uncertainty on the state space we can show that every set $A \in \mathcal{B}(\mathcal{S})$ can be reached from every $s_0 \in \mathcal{S}$ in a given number of steps m with positive probability. Consider a disc of radius R_s, enclosing the $\mathcal{X}$ part of the state space, where R_s is the maximal distance of a point to the origin (exemplarily illustrated in Figure 1). Note that for the derivation of maximal number of steps m, the angular part of $\mathcal{S}$ does not need to be considered since we only have dynamic constraints on the positions and $\vartheta(x)$ can always be covered in one step.

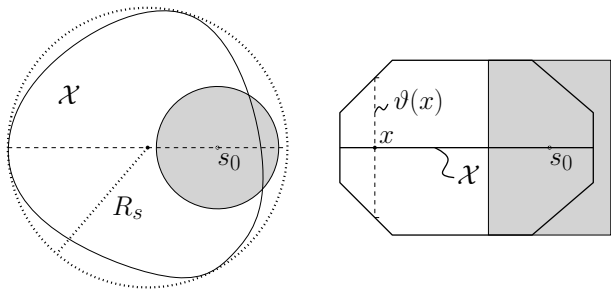

Fig. 1: *Left:* Top down view on $\mathcal{X}$ and the density support after the k-th step starting from s_0. *Right:* The angular plain of $\mathcal{S}$ for the dashed line on the left.

First consider the states which are not affected by the border, i.e., the proposal range $\Theta(s_k)$ is not constrained. The density support will always be a volume in $\mathcal{S}$ but for the sake of visualization we mainly depict the $\mathcal{X}$-part in Figure 2 for the first three steps. Given any initial state s_0 there is a

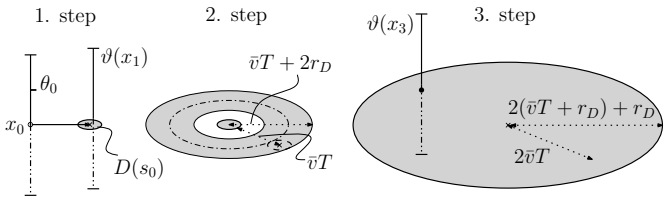

Fig. 2: Growing density support for 3 steps starting from s_0. The vertical bars indicate the density in the angular dimension at some example states s_i. For the initial step only θ_0 is covered, while in the next steps the whole range is proposed.

uniform density on the disc $D(s_0) \subset \mathcal{X}$ and uniform density over the whole range of angles $\vartheta(x_1) = [-\pi, \pi]$, creating a cylinder in $\mathcal{S}$ after the first transition. The density after the second transition has again the shape of a cylinder in $\mathcal{S}$ but with a void since the step length is fixed with bounded noise. The void in the density closes with the third transition leading to a full cylinder of radius $R_3 = 2(\bar{v}T + r_D) + r_D$, where $\bar{v}T$ is the fixed step size according to (1). This can easily be derived by superposition of the possible movements from all states with non-zeros density as depicted in Figure 3.

From the third transition on, the support has no voids and it's diameter will grow on $\mathcal{X}$ with the rate

$$
\frac{\bar{v}T + r_D}{\text{transition}}. \qquad (9)
$$

Note again, that the angular space is always covered in one transition. Now, the maximal distance between any initial starting point and any set $A \in \mathcal{B}(\mathcal{S})$ is $2R_s$. Hence, using the growth-rate of the support (9), we can immediately bound the maximal number of steps to cover $\mathcal{S}$ by

$$
m \leq \left\lceil 3 + \frac{2R_s - 2(\bar{v}T + r_D) - r_D}{\bar{v}T + r_D} \right\rceil = \left\lceil \frac{2R_s + \bar{v}T}{\bar{v}T + r_D} \right\rceil. \qquad (10)
$$

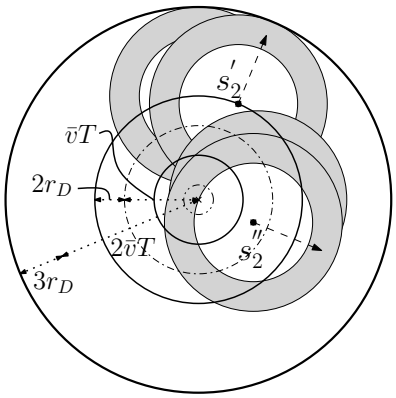

Fig. 3: Overlay of the density support after the second step and the transition densities starting from 4 example states. The two inner solid lines indicate the density support after the second step and the outer line the density after the third step. The gray area depicts the transition densities for 4 example transitions. Since there is a continuous density on the inner ring, it follows that the outer ring is fully covered.

Here, the initial three steps to close the void in the support are added, hence the distance covered by these steps is subtracted from the maximal distance.

In a next step we have to study the effect of the boundary on the growth of the support. For this purpose we introduce the auxiliary concept of an "unconstrained support", i.e., the support that would grow without the influence of the boundary. After some transition the unconstrained density support will intersect with the complement of $\mathcal{S}$ as already indicated in Figure 1. However, the proposal mechanism is designed to keep the chain inside the state space for all times and hence these overlapping densities will be "mapped back" into $\mathcal{S}$.

It remains to show that no gaps in the support occur for states where $\Theta(s_k)$ is restricted by the border. By design, the proposal orientations are restricted in way that the noise could bring the state onto the border. Therefore the support in Figures 2 and 3 will be truncated but still closed. This is shown in Figure 4 for a state on the border, which has the most restricted set $\Theta(s_k)$.

It has to be noticed that a small indent of the support occurs where the boundary and unconstrained support intersect (see the shaded areas in Figure 4). This depends on the angle in which boundary and unconstrained support intersect as well as on the curvature of the boundary and the bound on the noise. Nevertheless this gap will be "pushed" along the boundary with every transition. Due to the convexity of $\mathcal{X}$ and the upper limit on the curvature given in Assumption 1 the density support will merge from the left and the right after at least the maximal number of steps. The growth of the support for initial conditions near the boundary is only different in the sense, that the density support after the third transition will be truncated by the boundary. There will be no voids as shown in Figure 4 and the small indents that arises again, will be covered as discussed before. Therefore we can conclude that any state inside $\mathcal{S}$ that is covered by the "unconstrained" density support is also covered by the truncated support.

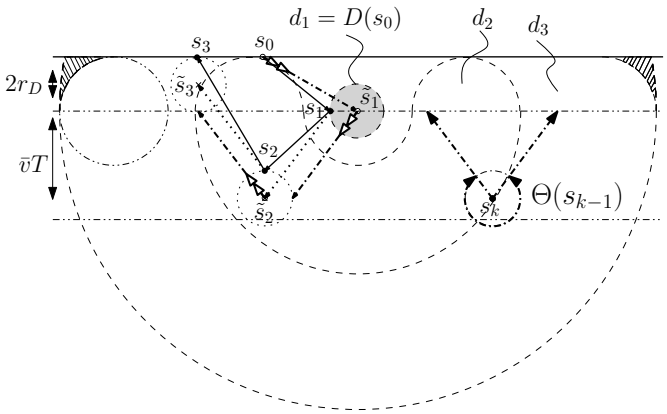

Fig. 4: The dashed lines depict the supports of the uncertainty for the k-th transition, labeled by d_k. For the example trajectory starting on the border at s_0 the dashed-dotted lines with double-arrows symbolize the accepted angles at the positions $\tilde{s}$. The dotted lines indicate the estimated movement and the solid lines the real movement under noise. For some state s_k on the right the orientations which can be proposed at this position are depicted and the circle on the left visualizes the curvature of the shaded indent.

Finally, the transition kernel given in Equation (5) consists of two terms where the first one is the *acceptance term*, which we will denote by $\mathcal{P}_{ac}$:

$$\mathcal{P}_{ac}(s_k, ds) = \frac{\mathbb{I}_{\{x \in D(s_k)\}}}{\gamma_{\mathcal{X}}(s_k)} \frac{\mathbb{I}_{\{\theta \in \Theta(s_k)\}}}{\gamma_{\Theta}(s_k)} \alpha(s_k) ds. \tag{11}$$

$\mathcal{P}_{ac}$ is the absolutely continuous part of the transition density with respect to the Lebesgue measure μ_{Leb} and under Assumption 4, with $\alpha > 0$, there is always a small probability contribution of this term. We showed that after m steps we cover every state in $\mathcal{S}$, where the states furthest away can only be covered after the last transition. This means that the density $\varrho(s)$ over the whole space has at least a small uniform part everywhere given by

$$\begin{aligned} \varrho(s) \geq \prod_{k=0}^{m} p(s_k, s) &= \prod_{k=0}^{m} \frac{\mathbb{I}_{\{x \in D(s_k)\}}}{\underbrace{\gamma_{\mathcal{X}}(s_k)}_{\text{constant} \neq 0}} \frac{\mathbb{I}_{\{\theta \in \Theta(s_k)\}}}{\underbrace{\gamma_{\Theta}(s_k)}_{\in (0, 2\pi]}} \alpha(s_k) \\ &\geq c^m \prod_{k=0}^{m} \alpha(s_k) \geq c^m \epsilon^m > 0 \end{aligned} \tag{12}$$

with $1 > c > 0$. Now the probability to reach any set $A \in \mathcal{B}(\mathcal{S})$ from any initial condition $s_0 \in \mathcal{S}$ is given by:

$$P^m(s_0, A) \geq \int_A \varrho(s) ds \geq c^m \epsilon^m \mu_{Leb}(A) > 0 \tag{13}$$

The m-th step transition probability from every $s_0 \in \mathcal{S}$ to any $A \in \mathcal{B}(\mathcal{S})$ is strictly positive, hence the Markov chain is μ_{Leb}-irreducible.

■

VI. CONVERGENCE AND ERGODICITY

Without any further assumptions we can obtain the following to convergence results for the Markov chain.

Proposition 6: Suppose that the Assumptions 1, 2 and 4 hold, then the Markov chain $\{s_k\}_{k\geq 0}$ has a unique stationary distribution π and further

(i) the chain is uniform ergodic and
(ii) there exists a measure $\nu_m(\mathcal{S})$ for some fixed m such that the chain converges to its stationary distribution π with the rate $2\rho^{\lfloor \frac{n}{m} \rfloor}$, where $\rho = 1 - \nu_m(\mathcal{S})$.

Proof: The existence and uniqueness of π is given by Proposition 3 and Proposition 5. In view of [10, Theorem 16.0.2] obtaining uniform ergodicity for the Markov chain is equivalent to prove that the state space $\mathcal{S}$ is a ν_m*-small set*. By the definition of ν_m-small sets in [10] we need to show that there exists a $m > 0$ and a nontrivial measure ν_m on $\mathcal{B}(\mathcal{S})$, such that

$$P^m(x, B) \geq \nu_m(B) \tag{14}$$

holds for all $s \in \mathcal{S}$, $B \in \mathcal{B}(\mathcal{S})$. Using the probability bound established in (13), Equation (14) immediately holds for $\mathcal{S}$, since we can identify the non-trivial measure to be $\nu_m(B) = c^m \epsilon^m \mu_{Leb}(B)$ and m to be the maximal number of steps derived in (10). The convergence rate can be immediately obtain from [10, Theorem 16.0.2]. Hence, for any x we have

$$||P^n(x, \cdot) - \pi|| \leq 2\rho^{\lfloor \frac{n}{m} \rfloor}, \tag{15}$$

where $||\cdot||$ denotes the *total variation norm* as defined in [10, Chap. 13]. ■

VII. CASE STUDY

We use the same scenario as in [7] to verify the results. Consider a a circular search space $\mathcal{X}$ around the origin with radius $R = 400\text{m}$ and the concentration field given by

$$C(x) = \sum_{i=1}^{3} w_i e^{-m\|x - x_{s,i}\|}, \tag{16}$$

with three maxima at $x_{s,1} = (-80, 50)$, $w_1 = 0.3$, $x_{s,2} = (0, -100)$, $w_2 = 0.7$, $x_{s,3} = (140, 120)$, $w_3 = 0.9$ and parameter $m = 0.1$.

The speed of the AUV is fixed to $\bar{v} = 1.5\frac{\text{m}}{\text{s}}$, the sampling time of the system is $T = 20\text{s}$ and the choice of the acceptance criterion α is the same as in [7]

$$\alpha(s_k) = 1 - e^{-(k[C(x_k)+\delta]^J)}\lambda(s_k). \tag{17}$$

Note that the border function $\lambda \in [0, 1]$ is used to ensure condition 3) in III, i.e., prevents the vehicles from leaving $\mathcal{X}$. The detailed construction of λ can be found in [7], where the thresholds have to be slightly modified to account for the new acceptance boundary $\mathcal{X}_n$. For $\delta = 7.5 \cdot 10^{-8}$ and the tuning parameters fixed at $k = 100$, $J = 0.7$, the minimal acceptance probability is given by

$$\epsilon = 1 - e^{-(k\delta)^J} = 0.01 \tag{18}$$

The probability bound ϵ is mainly introduced for the derivation of the convergence rate in (15), since the values of (16) will essentially never be equal to zero. One can assume that this holds also in practice as a concentration value of zero will possibly not appear in fluids and hence a small acceptance probability always exists.

A. Numerical Results

The distributions of the vehicle positions are used to localize the sources, where regions with higher density indicate their locations. Therefore, we study the effect of different bounds r_D for the noise on the shape of the numerically converged distributions, which can be treated as stationary, according to the theoretical results. Further we analyze the convergence rates under presence of noise. The error between two distribution is measured by the total variation distance (TV-distance), defined as half the L^1-norm on a countable state space [13].

Bound on the noise: The noise w which is added to the movement of the vehicles reflects disturbances which push the vehicle away from its track and position. We assume that the bound r_D is small compared to the step size, but for robustness analysis we considered noise levels of $r_D = \{\frac{1}{15}\bar{v}T, \frac{1}{3}\bar{v}T, \frac{2}{3}\bar{v}T\}$. For larger noise levels the distance between boundary of $\mathcal{X}$ and $\mathcal{X}_n$ becomes larger, creating a broader zone in which the vehicles are forced to turn. This can be observed for the extreme case of $r_D = \frac{2}{3}\bar{v}T$ in Figure 5, where the density increases on a circle whose radius becomes smaller for increasing noise levels. The annulus between this circle and the outer boundary has a low density since it is only visited by the vehicles if they are taken there by noise.

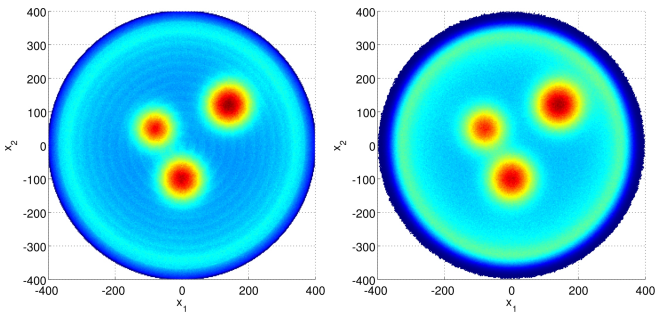

Fig. 5: Top-down view of the discretized spatial distributions of 5 vehicles and 100000 hours run time. The red areas indicate a high density. The space $\mathcal{X}$ was gridded with 300 bins in each dimension. *Left*: noised bound $r_D = 0\,\text{m}$, *Right*: $r_D = 20\,\text{m}$

Convergence rate: The convergence rate ρ from Equation (15) is given as

$$\rho = 1 - \nu_m(\mathcal{S}) = 1 - c^m \epsilon^m \mu_{Leb}(\mathcal{S}). \tag{19}$$

Given the circular search space $\mathcal{X}$ the state space $\mathcal{S}$ is composed of a cylinder of height $\vartheta_{min} = \min_{x \in \mathcal{X}} \vartheta(x)$ and two truncated cones of height $h = \frac{2\pi - \vartheta_{min}}{2}$. Selecting the noise level to be $r_D = 10\,\text{m}$, the minimal range of angles becomes $\vartheta_{min} = 2\arccos\left(\frac{(\bar{v}T)^2 + R^2 - (R - r_D)^2}{2\bar{v}TR}\right) \approx 2.391$.

The Lebesgue measure in (19) is then given by $\mu_{Leb}(\mathcal{S}) \approx 2.969 \cdot 10^6$. We obtain the maximal number of steps $m = 21$ from Equation (10) with the parameter chosen before and the radius $R_s = R$. With the normalization factors $\gamma_{\mathcal{X}} \approx 314$ and $\gamma_{\Theta} = 2\pi$ and ϵ from (18) the convergence rate is $\rho = 1 - \left(\frac{1}{\gamma_{\mathcal{X}}}\frac{1}{\gamma_{\Theta}}\right)^m \epsilon^m \mu_{Leb}(\mathcal{S}) \approx 1 - 3.384 \cdot 10^{-126}$. However, this rate is only of theoretical relevance, since it is bounded in a very conservative way.

The numerical results suggest a much faster convergence rate in practice. We can describe the linear behavior in the double-logarithmic plot of Figure 6 by

$$\log ||P^n(x,\cdot) - \pi|| \approx \log d - p \cdot \log n, \qquad (20)$$

where d is some constant. For the linear slope we obtain a value of $p \approx 0.5$ which corresponds to a sublinear convergence rate $||P^n(x,\cdot) - \pi|| \approx \frac{d}{\sqrt{n}}$, which is independent of the noise level (see Figure 7).

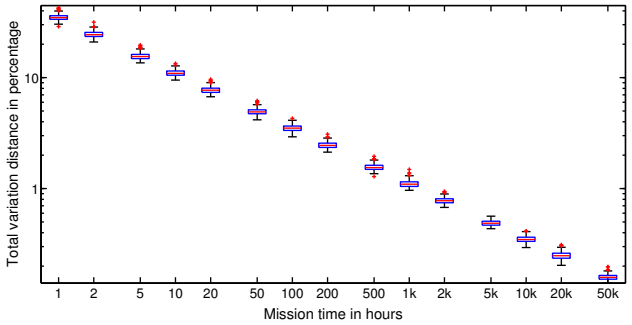

Fig. 6: TV-distances (in %) of distributions for 5 AUVs with respect to a distribution of 5 AUVs and 10^6 hours. The red marks inside the boxes are the medians of each group. The edges show the 25th and 75th percentiles while outliers are plotted individually. For each group, 200 independent MC runs were simulated. The noise level is $r_D = 3\,\mathrm{m}$

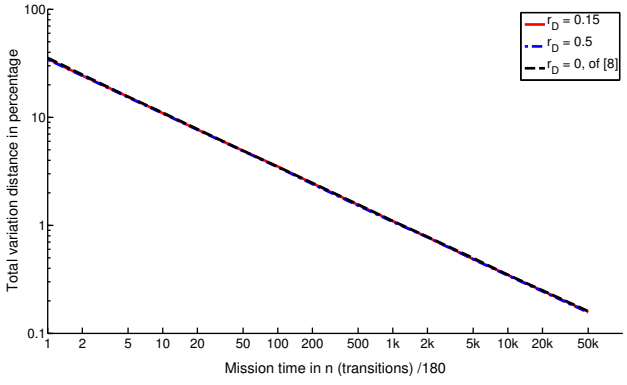

Fig. 7: Comparison of convergence rates in TV-distance (in %) for different noise levels. Depicted are the medians from 200 independent MC runs of distributions for 5 AUVs with respect to a distribution of 5 AUVs and $n = 10^6 \cdot 180$

VIII. CONCLUSIONS AND FUTURE WORKS

We presented an extension to our previously developed method for localization of sources, applicable to search scenarios for autonomous vehicles. The existence results of an invariant distribution hold in the presence of noise. We provided convergence guarantees for the MCMC-algorithm, by showing that the algorithm converges uniformly to a unique stationary distribution. Having established the theoretical foundations, a performance comparison to other approaches can be conducted next to motivate the practical use. The numerical results suggest a sublinear convergence rate with the square root of the number of steps. Investigation on this convergence rates will be subject to further research. Preliminary results indicate that the algorithm performs reliable with the same rate under measurement noise. Future work will mainly involve the improvement of the convergence, by including a cooperation scheme based on communications among the vehicles as well as different proposal schemes.

ACKNOWLEDGEMENT

The authors would like to thank Debasish Chatterjee and Andreas Milias-Argeitis for their insight and valuable comments.

REFERENCES

[1] Bachmayer, R. and Leonard, N.E. Vehicle networks for gradient descent in a sampled environment. In *Proceedings of the 41st IEEE Conference on Decision and Control, 2002*, volume vol.1, pages 112 – 117, Mar. 2003.

[2] E. Burian, D. Yoerger, A. Bradley, and H. Singh. Gradient search with autonomous underwater vehicles using scalar measurements. In *Autonomous Underwater Vehicle Technology, 1996. AUV '96., Proceedings of the 1996 Symposium on*, pages 86–98, 1996.

[3] J. Cochran and M. Krstic. Nonholonomic source seeking with tuning of angular velocity. *Automatic Control, IEEE Transactions on*, 54(4):717–731, 2009.

[4] M. P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice Hall, Mar. 1976.

[5] J. Farrell, S. Pang, and W. Li. Plume mapping via hidden markov methods. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 33(6):850–863, 2003.

[6] O. Hernández-Lerma and J. B. Lasserre. *Markov Chains and Invariant Probabilities*. Birkhäuser Basel, 1 edition, Apr. 2003.

[7] S. Huck, P. Hokayem, D. Chatterjee, and J. Lygeros. Stochastic localization of sources using autonomous underwater vehicles. In *American Control Conference (ACC), 2012*, pages 4192–4197, 2012.

[8] A. Martins, J. M. Almeida, E. Silva, and F. L. Pereira. Hybrid maneuver for gradient search with multiple coordinated AUVs. In *Proceedings in 5t IFAC Symposium on Intelligent Autonomous Vehicles (IAV 2004 Lisaboa July 2004)*, 2004.

[9] A. Mesquita, J. Hespanha, and K. Åström. Optimotaxis: A stochastic multi-agent optimization procedure with point measurements. In *Hybrid Systems: Computation and Control*, volume Volume 4981/2008 of *Lecture Notes in Computer Science*, pages 358–371. Springer Berlin / Heidelberg, 2008.

[10] S. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Cambridge University Press, 2 edition, Apr. 2009.

[11] B. Moore and C. Canudas-de Wit. Source seeking via collaborative measurements by a circular formation of agents. In *American Control Conference (ACC), 2010*, pages 6417–6422, 2010.

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

[13] J. S. Rosenthal. Convergence rates for markov chains. *SIAM Review*, 37:387, 1995.

[14] H. Royden. *Real Analysis*. Prentice Hall, 3 edition, Feb. 1988.

[15] C. Zhang, D. Arnold, N. Ghods, A. Siranosian, and M. Krstic. Source seeking with non-holonomic unicycle without position measurement and with tuning of forward velocity. *Systems & Control Letters*, 56(3):245–252, Mar. 2007.