



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Khong, SZ;Tan, Y;Manzie, C;Nešić, D

Title:

Multi-agent source seeking via discrete-time extremum seeking control

Date:

2014-09-01

Citation:

Khong, S. Z., Tan, Y., Manzie, C. & Nešić, D. (2014). Multi-agent source seeking via discrete-time extremum seeking control. *Automatica*, 50 (9), pp.2312-2320. <https://doi.org/10.1016/j.automatica.2014.06.009>.

Persistent Link:

<https://hdl.handle.net/11343/299567>

Multi-agent source seeking via discrete-time extremum seeking control [★]

Sei Zhen Khong ^{a,1}, Ying Tan ^b, Chris Manzie ^c, Dragan Nešić ^b

^a*Department of Automatic Control, Lund University, SE-221 00 Lund, Sweden*

^b*Department of Electrical and Electronic Engineering, The University of Melbourne, Parkville, VIC 3010, Australia*

^c*Department of Mechanical Engineering, The University of Melbourne, Parkville, VIC 3010, Australia*

Abstract

Recent developments in extremum seeking theory have established a general framework for the methodology, although the specific implementations, particularly in the context of multi-agent systems, have not been demonstrated. In this work, a group of sensor-enabled vehicles is used in the context of the extremum seeking problem using both local and global optimisation algorithms to locate the extremum of an unknown scalar field distribution. For the former, the extremum seeker exploits estimates of gradients of the field from local dithering sensor measurements collected by the mobile agents. It is assumed that distributed coordination which ensures uniform asymptotic stability with respect to a prescribed formation of the agents is employed. An inherent advantage of the frameworks is that a broad range of nonlinear programming algorithms can be combined with a wide class of cooperative control laws to perform extreme source seeking. Semi-global practical asymptotically stable convergence to local extrema is established in the presence of field sampling noise. Subsequently, global extremum seeking with multiple agents is investigated and shown to give rise to robust practical convergence whose speed can be improved via computational parallelism. Nonconvex field distributions with local extrema can be accommodated within this global framework.

Key words: Extremum seeking, local and global optimisation, multi-agent systems, cooperative control

1 Introduction

The problem of localising the source of an unknown or uncertain scalar field environment has attracted considerable research attention. Such a problem arises from the need to localise a source flow (e.g. chemical pollution) in the ocean with a fleet Autonomous Underwater Vehicles (AUVs) (Fiorelli et al., 2006; Leonard et al., 2010), for example. The signal or field distribution can be the concentration of a chemical/biological/acoustic/electromagnetic entity, whose strength decays away from the source. Thus, an extreme

source seeking which relies on gradient information is also called gradient climbing/descending in the literature. The lack of knowledge about the scalar signal field to be optimised suggests that extremum seeking would be a suitable tool for tackling the problem.

In Zhang et al. (2007); Cochran and Krstic (2009), the case of a single autonomous vehicle without position measurements is investigated. A hybrid controller is developed in Mayhew et al. (2007, 2008) for a single-vehicle extreme source seeking using successive line minimisations with directional changes based on conjugate vectors. Non-local practical stability results are obtained for a certain class of signal strength distributions, with a margin of robustness to measurement noise. It is worth noting that the proposed hybrid controllers do not exploit gradient approximations in their operations.

There are potential advantages to using multiple agents for extremum seeking instead of a single one such as robustness to vehicle failure, scalability, increased reliability and search speed etc. In Ögren et al. (2004), a net-

[★] This work was supported by the Swedish Research Council through the LCCC Linnaeus centre and the Australian Research Council. This paper was not presented at any IFAC meeting.

Email addresses: seizhen@control.lth.se (Sei Zhen Khong), yingt@unimelb.edu.au (Ying Tan), manziec@unimelb.edu.au (Chris Manzie), dnesic@unimelb.edu.au (Dragan Nešić).

¹ Tel. +46 462229743. Fax +46 46138118.

work of sensor-enabled agents is employed to seek out local extrema in a distributed environment. Collectively the mobile sensors form an intelligently interacting sensor array and they are coordinated using virtual bodies and artificial potentials. Continuous-time gradient descent updates are applied to selected virtual leaders for the network to cooperatively perform the gradient climbing task. In Biyik and Arcaç (2008), a control algorithm is developed for a vehicle to lead a group of others in a prescribed formation to the source via a passivity-based distributed coordination framework and discrete-time Newton method. There, the sensor-enabled leader is driven by a reference velocity returned by the extremum seeker while the sensor-disabled followers reconstruct this information adaptively (Bai et al., 2011).

This paper considers a fleet of autonomous point-mass vehicles endowed with uniformly asymptotically stabilising cooperative control laws. It is assumed that only samples of measurements of the field are available, i.e. a vehicle cannot measure a continuum of the signal field. Such an assumption is justified, for example, when data collection is costly and/or time-consuming. Extremum seeking is performed in the discrete time within the general frameworks of Khong et al. (2013). While all the aforementioned approaches independently consider the closed-loop behaviour of the system, the framework methodology developed in Khong et al. (2013) enables the convergence of the overall system to be addressed simply by demonstrating the proposed algorithms satisfy certain conditions. Using gradient-based extremum seeking, local dither measurements are taken by a group of mobile sensors and then used to estimate the gradients of the objective field distribution. This information is utilised by the discrete-time extremum seeking controller to determine the subsequent fleet manoeuvre. The fleet formation is designed in such a way that the measurements required to approximate the gradients can be taken simultaneously, i.e. the dithering motions of the leader in Biyik and Arcaç (2008) are avoided altogether. A broad class of optimisation algorithms fit within the framework, allowing the user to base the selection on the complexity of implementation, speed of convergence, robustness, etc. at the control design stage. Furthermore, the fleet formation can be maintained by applying various consensus algorithms known in the literature (Ren and Beard, 2008; Bai et al., 2011). Semi-global practical asymptotic stability of local extrema is established in the presence of norm-bounded noise to the sampled measurements on the objective field distribution. The main contribution of the paper is that of identifying generic conditions about an optimisation algorithm and a cooperative control law such that combining the two yields a convergent extremum seeking scheme which manoeuvres a group of agents towards the extreme source.

The aforementioned extremum seeking framework exploits gradient information in their control updates. As such, the agents governed by such control laws would

only converge to local extrema, unless the underlying field distribution is convex. By assuming that the scalar field is defined on a compact subset, this paper proposes a multi-vehicle global extremum seeking framework, which accommodates several sampling-based non-convex optimisation algorithms such as Shubert and DIRECT (Jones et al., 1993). These algorithms place emphasis on both local and global search of an extremum. They collect scattered samples with the vehicles in the search space in succession following a set of sensibly designed rules which guarantee practical convergence in the face of bounded measurement noise. It is important to note that Khong et al. (2013) also enables global optimisation techniques to be deployed, although their specific implementation in the context of multi-agent source seeking is not discussed. The global extremum seeking framework in this paper takes departure from the multi-agent literature where a formation of the vehicles is prescribed, as inspired by such biological problems as flocking, schooling, and swarming. These behaviours often stem from the lack of GPS, communication constraints, and limited processing powers of the animals. In contrast, this paper assumes that each of the agent is equipped with the ability to implement the global extremum seeking controllers and loosens the constraint on formational maintenance. Two vehicle routing algorithms are proposed, taking into consideration the restricted communication ranges between the agents. Parallel measurements by multiple vehicles is exploited to accelerate the search for global extreme sources. Note that other communication issues such as dropouts and time-varying topology are not considered as they add another level of robustness analysis which would distract the reader from the main focus of extremum seeking.

The paper evolves along the following lines. First, a unifying discrete-time gradient-based extremum seeking control framework from Khong et al. (2013) is reviewed in the next section. In Section 3. it is shown how the framework can be adapted for extreme source seeking with a single autonomous vehicle. Source seeking with multiple sensor-enabled agents is considered in Section 4. Global extremum seeking is examined in Section 5. Simulation examples are provided in Section 6, followed by conclusions in Section 7.

2 Discrete-time extremum seeking

The real and natural numbers are denoted $\mathbb{R}$ and $\mathbb{N}$ respectively. Given a vector x in $\mathbb{R}^n$, its components are denoted by x^i for $i = 1, 2, \dots, n$. A function $\gamma : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is of class- $\mathcal{K}$ (denoted $\gamma \in \mathcal{K}$) if it is continuous, strictly increasing, and $\gamma(0) = 0$. If γ is also unbounded, then $\gamma \in \mathcal{K}_{\infty}$. A continuous function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is of class- $\mathcal{KL}$ if for each fixed t , $\beta(\cdot, t) \in \mathcal{K}$ and for each fixed s , $\beta(s, \cdot)$ is decreasing to zero (Khalil, 2002). The Euclidean norm is denoted $\|\cdot\|_2$. Let $\mathcal{X}$ be a Banach space with norm $\|\cdot\|$. Given any subset $\mathcal{Y}$ of $\mathcal{X}$

and a point $x \in \mathcal{X}$, define the distance of x from $\mathcal{Y}$ as $\|x\|_{\mathcal{Y}} := \inf_{a \in \mathcal{Y}} \|x - a\|$.

Consider the optimisation problem:

$$y^* := \max_{x \in \Omega} Q(x), \quad (1)$$

where $Q : \Omega \subset \mathbb{R}^n \rightarrow \mathbb{R}$ is an unknown, stationary (time-invariant), and Lipschitz continuous function or scalar field distribution which takes its maximum value on $\mathcal{C} \subset \Omega$, i.e. $Q(x) = y^*$ for all $x \in \mathcal{C}$. It is assumed that Q can only be sampled discretely in its domain Ω . Let Σ be a discrete-time extremum seeking algorithm for (1).

In the presence of bounded additive perturbations on the measurements as illustrated in Figure 1, i.e. $y_k = Q(x_k) + w_k$ with $|w_k| \leq \nu$ for $k = 0, 1, \dots$ and some $\nu > 0$, the following assumption is important to establish convergence of the extremum seeking schemes in subsequent sections.




Fig. 1. Extremum seeking algorithm with noisy output measurement.

Assumption 1 *The extremum seeking controller Σ satisfies the following conditions:*

- (i) Σ is time-invariant. Denote by $\{\hat{x}_k\}_{k=0}^\infty \subset \Omega$ the output sequence Σ generates based on input to Σ , $\{\hat{y}_k\}_{k=0}^\infty$, where $\hat{y}_k := Q(\hat{x}_k)$ denotes the ideal measurements. Σ is causal in the sense that the output at any time $N \in \mathbb{N}$, i.e. $\hat{x}_N$, is determined based only on $\hat{x}_k$ and $\hat{y}_k$ for $k = 0, 1, \dots, N-1$, that is the past probe values to Q and the corresponding measurements.
- (ii) Denote by $\mathcal{S}(\hat{x}_0)$ the set of all admissible output sequences of Σ with respect to the initial point $\hat{x}_0$. There exists a class- $\mathcal{KL}$ function β such that for any initial point $\hat{x}_0 \in \Omega$, all outputs $\hat{x} \in \mathcal{S}(\hat{x}_0)$ satisfy for some $\delta \geq 0$,

$$\|\hat{x}_k\|_{\mathcal{C}} \leq \beta(\|\hat{x}_0\|_{\mathcal{C}}, k) + \delta \quad \forall k \geq 0. \quad (2)$$

- (iii) Let $y_k := Q(x_k) + w_k$, where $w_k \in \mathbb{R}$. Denote by $\{x_k\}_{k=0}^\infty$ the output sequence Σ generates based on the perturbed input $\{y_k\}_{k=0}^\infty$. The pair (x, y) is multi-step consistent/close (Nešić et al., 1999) with $(\hat{x}, \hat{y})$, in the sense that for any positive (Δ, η) and $N \in \mathbb{N}$, there exists a $\nu > 0$ such that if $\|x_0\|_{\mathcal{C}} \leq \Delta$ and $|w_k| \leq \nu$ for $k = 0, \dots, N$, then there exists a $\hat{x} \in$

$\mathcal{S}(x_0)$ satisfying

$$\|x_k - \hat{x}_k\|_2 \leq \eta \quad \text{for } k = 0, 1, \dots, N.$$

Assumption 1 covers a wide class of optimisation algorithms. The next subsection provides some examples.

2.1 Examples of extremum seeking algorithms

It is often the case that gradient-based extremum seeking algorithms can be serially decomposed into a derivative estimator and a nonlinear programming method as shown in Figure 2.




Fig. 2. A gradient-based extremum seeking controller paradigm.

Procedure 1 *Let the initial output of the extremum seeking controller be x_0 . As determined by the derivative estimator, the following length- p sequence of points can be used to probe the field distribution Q along the directions given by the basis vectors $z_1, \dots, z_m$:*

$$(x_0 + d_1(x_0), \dots, x_0 + d_p(x_0)), \quad (3)$$

where $d_i : \Omega \rightarrow \mathbb{R}^n$ denote the dither signals. The corresponding outputs of Q are then sampled and collected by the derivative estimator to numerically approximate the first N -order partial derivatives of Q at x_0 using the Euler methods, trapezoidal rule, or the more sophisticated Runge-Kutta method, as needed by the optimisation algorithm. Exploiting this information, the optimisation algorithm can then update its output to x_1 , and the series of steps described above repeats.

Examples of extremum seeking algorithms satisfying Assumption 1 are the gradient descent and Newton methods. Teel and Popović (2001) describes a general class of such algorithms modelled by difference inclusions; see (Khong et al., 2013, Section 6) for more details.

3 Single-agent extreme source seeking

This section demonstrates how extreme source seeking of a $Q : \Omega \subset \mathbb{R}^n \rightarrow \mathbb{R}$ described in the previous section can be accomplished with a single controllable point-mass vehicle. It is assumed that measurements are corrupted by norm-bounded noise. Here, $n = 1, 2$, or 3 . The class of controllable vehicles include, for example, autonomous vehicles considered in Ögren et al. (2004); Biyik and Arcaç (2008). Semi-global practical asymptotic stability of the extrema of field distributions is established.

Assumption 2 The vehicle dynamics are

$$\dot{x} = f(\xi, u, v); \quad \dot{\xi} = g(\xi, u), \quad (4)$$

where $x(t) \in \mathbb{R}^n$ and $\xi(t) \in \mathbb{R}^p$ denote, respectively, the position with respect to an inertial frame and internal dynamics of the vehicle, $u(t) \in \mathbb{R}^m$ the (cooperative) control input, and $v(t) \in \mathbb{R}^n$ the reference velocity. Both f and g are locally Lipschitz in each argument. The vehicle is controllable in the following sense. Suppose $x(0) = x_0$ for some $x_0 \in \mathbb{R}^n$. Given any sequence $\{x_k\}_{k=0}^\infty$ in $\mathbb{R}^n$, there exist an increasing sequence $\{\tau_k\}_{k=1}^\infty$ in $\mathbb{R}$ and a piecewise continuous reference velocity v such that when applied to (4) with $u = 0$ results in $x(\tau_k) = x_k$ for $k = 1, 2, \dots$. Suppose also that given any $\epsilon > 0$, there exists a $T > 0$ such that $\|v(t)\|_2 \leq \epsilon$ for $t \in [\tau_k, \tau_{k+1}]$ if $\tau_{k+1} - \tau_k \geq T$.

The last bit of the assumption means that a vehicle travelling at a lower speed takes longer time to get from one point to another.

Remark 3 The vehicle dynamics in Assumption 2 may also take the form

$$\ddot{x} = f(\xi, u, v); \quad \dot{\xi} = g(\xi, u),$$

where $v(t) \in \mathbb{R}^n$ now denotes the reference applied force, which serves the same controllability function as above.

Various types of vehicular dynamics satisfy Assumption 2. The controllability of rigid body attitude dynamics is studied in Reyhanoglu et al. (1999). The single and double integrator dynamics considered in the next example are well-studied in the multi-agent literature (Ren and Beard, 2008; Bai et al., 2011).

Example 4 An example which satisfies a stronger version of the above assumption is an autonomous vehicle with dynamics $\ddot{x} = u$, where u denotes the control force and has the same dimension as x . This model may arise from applying output feedback linearisation to the dynamics of a wheeled mobile robot about the hand position (Lawton et al., 2003); see Section 4.2 for more details. The following argument generalises the velocity assignment in (Biyik and Arcak, 2008, Section II). Suppose $x(0) = x_0 \in \mathbb{R}^n$ and for some $\tau_2 > \tau_1 > 0$, one needs to design a continuous reference velocity v to drive the vehicle to $x(\tau_1) = x_1$ and $x(\tau_2) = x_2$. This can be achieved with the reference velocity:

$$\dot{x}(t) := v(t) \quad \text{for } t \in [0, \tau_2],$$

where

$$v(t) := (x_1 - x_0) \frac{1}{\tau_1} \left(1 - \cos \left(\frac{2\pi}{\tau_1} t \right) \right) \quad \text{for } t \leq \tau_1$$

and

$$v(t) := (x_2 - x_1) \frac{1}{\tau_2 - \tau_1} \left(1 - \cos \left(\frac{2\pi}{\tau_2 - \tau_1} (t - \tau_1) \right) \right) \quad \text{for } t \in (\tau_1, \tau_2].$$

Integrating the expressions above yields the desired results. Observe that v is continuously differentiable and

$$(v(t), \dot{v}(t))|_{t=0, \tau_1, \tau_2} = (0, 0).$$

Moreover, $\|v(t)\|_2$ is inversely proportional to $\tau_{k+1} - \tau_k$ for $t \in [\tau_k, \tau_{k+1}]$. As such, given a sequence $\{x_k\}_{k=0}^\infty$ in $\mathbb{R}^n$ and an increasing sequence $\{\tau_k\}_{k=1}^\infty$, there always exists a continuously differentiable reference velocity v which when applied to the autonomous vehicle $\dot{x}(t) = v(t)$ with $x(0) = x_0$ yields that $x(\tau_k) = x_k$ for $k = 1, 2, \dots$ by following the above argument inductively. Note that this reference velocity is not unique. Realistically, τ_k 's would be sufficiently spaced so that the control force does not exceed the physical constraints of the vehicle.




Fig. 3. Source seeking with a single agent.

The main result of this section is stated next. Let $Q : \Omega \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be a Lipschitz continuous field distribution with the set of maximisers $\mathcal{C}$. Given a vehicle satisfying Assumption 2 and an extremum seeking controller satisfying Assumption 1, consider the extreme source seeking setup in Figure 3, where the vehicle samples Q at the user-defined time instants $\{\tau_k\}_{k=1}^\infty$. Let w_k be the additive measurement noise. By the preceding arguments, suppose a piecewise continuous reference velocity v has been designed so that the position trajectory of the vehicle satisfies $x(\tau_k) = x_k$ for $k = 1, 2, \dots$, where x_k is the output of the extremum seeking controller with its input being $y_k = Q(x_k) + w_k$, the noise-corrupted samples of Q .

Theorem 5 Let the extremum seeking configuration be as described above. Given any (Δ, μ) such that $\Delta, \mu > \delta$, where $\delta \geq 0$ is as in Assumption 1(ii), there exist $\nu > 0$ and a $\bar{\beta} \in \mathcal{KL}$ such that for any $\|x_0\|_{\mathcal{C}} \leq \Delta$, if $|w_k| \leq \nu$, then

$$\|x_k\|_{\mathcal{C}} \leq \bar{\beta}(\|x_0\|_{\mathcal{C}}, k) + \mu \quad (5)$$

for all $k = 0, 1, \dots$

PROOF. The proof can be established using similar arguments in (Khong et al., 2013, Thm. 19). It exploits the multi-step consistency and time-invariance of the extremum seeking algorithm. $\square$

Theorem 5 demonstrates that semi-global practical asymptotic convergence of the extremum seeking scheme can be achieved with a single controllable agent. In the succeeding section, a network of multiple agents is examined within the context of extremum seeking control.

4 Multi-agent extremum seeking

While the problem of extreme source seeking can be tackled using a single expensive mobile agent as in the previous section, there are cases where multiple sensor-enabled vehicles are needed to move slowly in a formation towards the source (Ögren et al., 2004). For instance, a fleet of economical AUV's may be deployed to locate a source in the ocean. By working collaboratively, they are less prone to failures due to vehicular malfunction as opposed to the single-agent case. This section considers the problem of steering a group of sensor-equipped agents to the source of the signal field using the extreme seeking framework described in Section 2, together with a generic set of distributed coordination rules which maintain the agents' formation. It is assumed that the extremum seeking controller has access to all the samples collected by the sensing agents. However, it does not function as a supervisory controller and the agents are coordinated in a distributed fashion (Bai et al., 2011; Ren and Beard, 2008). As in the single-agent section, semi-global practical asymptotic stability is established by exploiting the uniform asymptotic stability of the fleet formation.

The formation of the mobile sensors can be chosen in such a way that facilitates the estimation of the gradients of the environment field. For instance, using the simplest first-order gradient approximation method, the Euler's finite-difference, the gradient of a field distribution $Q : \mathbb{R}^2 \rightarrow \mathbb{R}$ at $x \in \mathbb{R}^2$ can be approximated with three samples x , $x + hz^1$, and $x + hz^2$, where z^1 and z^2 denote the canonical basis vectors for $\mathbb{R}^2$ and h is a small step size. In particular, the approximation is given by

$$\nabla Q(x) \approx \begin{bmatrix} \frac{Q(x+hz^1) - Q(x)}{h} \\ \frac{Q(x+hz^2) - Q(x)}{h} \end{bmatrix}. \quad (6)$$

In this case, a network of three agents can be deployed and their formation selected to be the vertices of a right-angled triangle. Likewise, in $\mathbb{R}^3$, four agents may be deployed to form a tetrahedron. If more complicated gradient approximation methods are employed such as the family of Runge-Kutta, more samples points would be

needed and the number and formation of the agents can be decided accordingly.

Let the topology of information exchange between N number of dynamic agents be modelled by a graph. A desired formation for the purpose of derivative estimation may be given by

$$\mathcal{P} = \{z^i \mid |z^i| = b^i; i = 1, 2, \dots, \ell\}, \quad (7)$$

where $b^i > 0$, z^i denotes the distance between two agents connected by a link, and ℓ the total number of links in the graph. Note that $\mathcal{P}$ can also be chosen to be position-based and/or updated over time; see Section 4.2 for an example. The dynamics of the agents are assumed to satisfy Assumption 2 and are given by

$$\dot{x}^i = f(\xi^i, u^i, v^i); \quad \dot{\xi}^i = g(\xi^i, u^i) \quad i = 1, 2, \dots, N.$$

The cooperative control u^i is a function of x^j and $\dot{x}^j$ if the j^{th} agent is linked to the i^{th} agent or in certain consensus control frameworks where virtual leaders are exploited, if the j^{th} agent is a virtual leader. It is nominally zero when the vehicle network achieved the desired formation. The following assumption prescribes an objective for a formation coordination framework to be used in conjunction with extremum seeking.

Assumption 6 *Given N number of mobile agents satisfying Assumption 2 and a desired formation, the distributed coordination control law is such that the formation is a uniformly asymptotically stable equilibrium.*

In the case where (7) is used, the assumption above means there exist a $c > 0$ and a $\beta \in \mathcal{KL}$ such that

$$\|z^i(t)\|_{\mathcal{P}} \leq \beta(\|z^i(0)\|_{\mathcal{P}}, t) \quad \forall \|z^i(0)\|_{\mathcal{P}} \leq c, i = 1, \dots, \ell.$$

Examples of consensus algorithms that satisfy the assumption above are provided in the next subsections. Extremum seeking with a group of vehicles is examined below.

Given a Lipschitz continuous $Q : \Omega \subset \mathbb{R}^n \rightarrow \mathbb{R}$ with Lipschitz constant $L > 0$, consider the gradient-based multi-agent extremum seeking scheme illustrated in Figure 4, where $x^1, \dots, x^N$ represent the state trajectory of agent 1 to N respectively and $d_1, \dots, d_N$ denote the dither functions as in (3). In particular, it is assumed that given any $x \in \mathbb{R}^n$, $x + d_1(x), \dots, x + d_N(x)$ define a formation (e.g. vertices of a triangle/tetrahedron) by which the derivatives of Q at x can be well-approximated with the information $Q(x + d_1(x)), \dots, Q(x + d_N(x))$. The dithers can be constant functions as in Biyik and Arcak (2008). On the other hand, Ögren et al. (2004) studies dither functions that vary with the positions of the vehicles and adapt their configuration in response to the




Fig. 4. Source seeking with multiple agents.

measurement noise to optimise gradient climbing. Reference velocities $v^1, \dots, v^N$ are assumed to be designed according to Assumption 2 to steer the vehicles' positions $x^1, \dots, x^N$ towards $x + d_1(x), \dots, x + d_N(x)$, similarly to the single-vehicle case. The formation of the vehicles is maintained with a consensus algorithm satisfying Assumption 6.

Because the agents are driven by both the reference velocity v and the cooperative control laws u , in general $x^i(\tau_k)$ will lie close to but not precisely on $x_k + d_i(x_k)$, by contrast with the single-agent case. Indeed, $x^i(\cdot)$ converges to $x_k + d_i(x_k)$ as time tends to infinity by the uniformly asymptotically stable property of the equilibria defining the desired formation. Nevertheless, this is sufficient to give rise to practical convergence to the extrema of the field distribution as demonstrated by the next result.

Theorem 7 *Suppose the extremum seeking controller in Figure 4 satisfies Assumption 1 and the agents with dynamics in Assumption 2 are coordinated with a control law satisfying Assumption 6 with respect to the formation defined by the dither functions. Also assume that the vehicles commence within the cooperative control law's region of attraction and are driven by appropriate reference velocities/forces towards destinations specified by the extremum seeking controller. Then given any (Δ, μ) such that $\Delta, \mu > \delta$, where $\delta \geq 0$ is as in Assumption 1(ii), there exist a $T > 0$ and a $\beta \in \mathcal{KL}$ such that for any $\|x_0\|_C \leq \Delta$, if $\tau_{i+1} - \tau_i \geq T \forall i = 0, 1, \dots$, then*

$$\|x_k\|_C \leq \beta(\|x_0\|_C, k) + \mu \quad (8)$$

for all $k = 0, 1, \dots$

PROOF. First note that the gradient motions of the vehicles as determined by the reference velocity need to be sufficiently slow so that the vehicles never leave the local region of attraction of the formation during

manoeuvres. By Assumption 2, this can be guaranteed by an $T_1 > 0$ such that $\tau_{i+1} - \tau_i \geq T_1 \forall i = 0, 1, \dots$

Now following the arguments in Theorem 5, there exists a $\nu > 0$ such that (8) holds if $|x_k + d_i(x_k) - x^i(\tau_k)| \leq \nu$ for all $i = 1, \dots, N$ and $k = 1, 2, \dots$. By the Lipschitz continuity of the environment distribution Q , the latter holds if $x^i(\tau_k)$ lies within a ball of radius $\frac{\nu}{L}$ centered at $x_k + d_i(x_k)$ for all i and k , where L is a Lipschitz bound for Q . In other words,

$$x^i(\tau_k) \in x_k + d_i(x_k) + \frac{\nu}{L} \bar{B}, \quad (9)$$

where $\bar{B}$ denotes the closed unit ball in $\mathbb{R}^n$. Since $x_k + d_i(x_k)$ is a uniformly asymptotically stable equilibrium for the i^{th} agent by Assumption 6, it follows that (9) can be ensured if the network is given enough time to manoeuvre towards its final formation during $[\tau_{k-1}, \tau_k]$. That is, there needs a sufficiently large gap $T_2 > 0$ between the time instants τ_{k-1} and τ_k . Defining $T := \max\{T_1, T_2\}$ completes the proof. $\square$

Theorem 7 identifies general conditions on an optimisation algorithm (Assumption 1) and a cooperative control law (Assumption 6), combining which guarantees convergence towards an extreme source of a field distribution for a class of controllable mobile agents satisfying Assumption 2. Examples of such optimisation algorithms in Section 2 and examples of such cooperative control laws in subsequent subsections demonstrate that a wide range of options are at the user's disposal. Various properties of these algorithms and control laws can be investigated and exploited at the design stage to decide which to implement on a given problem.

4.1 Examples of cooperative control laws

Assumption 6 can be achieved with a variety of cooperative control methods under different conditions; see Ren and Beard (2008), an example in which is provided in the next subsection. In Leonard and Fiorelli (2001); Ögren et al. (2004), the communication topology of the agents is modelled by a complete graph, i.e. there is an undirected (bidirectional) link between every two agents. Autonomous point-mass vehicles with fully actuated dynamics $\ddot{x} = u$ are considered. Virtual leaders and artificial potentials are introduced to manipulate the formation and motion of the network. In particular, collision avoidance can be achieved by appropriately setting the artificial potentials so that two vehicles 'repel' each other when they get too close. Local asymptotic stability of the equilibrium corresponding to the vehicles at rest at the global minimum of the sum of the artificial potentials is established using a kinetic and potential based Lyapunov function. Ögren et al. (2004) exerts motion on the virtual body to decouple the network manoeuvre from the

cooperative management of the network formation and regulates its speed to ensure uniform asymptotic stability. There, a continuous-time gradient descent method is exploited for the task of extremum seeking. Theorem 7 demonstrates that a rich class of discrete-time optimisation algorithms may be employed for the same purpose.

Biyik and Arcaç (2008) considers vehicles whose dynamics can be rendered passive via an internal feedback design and a communication structure given by an undirected graph. Only local neighbouring agents connected with an edge can exchange information. The reference velocity is additive to the vehicle's velocity. A systematic passivity-based cooperative control framework is proposed. Collision can be avoided by incorporating additional feedback terms in the control law. For position-based formation control, uniform asymptotic stability of a prescribed compact subset is established using passivity-based arguments. Also given is a characterisation of the region of attraction, which can at times be designed to be arbitrarily large. Biyik and Arcaç (2008) proposes a discrete-time Newton method to tackle the problem of extreme source seeking. This aspect is generalised by Theorem 7 to include optimisation algorithms modelled by difference inclusion (Teel, 2000), for instance.

4.2 An example of formation manoeuvre

This subsection examines the decentralised approach to formation manoeuvres of Lawton et al. (2003); Ren and Beard (2008) and demonstrates that it fits within the general framework of Theorem 7. It is also used as a basis of simulations in Section 6. Lawton et al. (2003) considers N number of wheeled robots with the following dynamics:

$$\begin{bmatrix} \dot{r}_x^i \\ \dot{r}_y^i \\ \dot{\theta}^i \\ \dot{v}^i \\ \dot{\omega}^i \end{bmatrix} = \begin{bmatrix} v^i \cos(\theta^i) \\ v^i \sin(\theta^i) \\ \omega_i \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \frac{1}{m^i} & 0 \\ 0 & \frac{1}{J^i} \end{bmatrix} \begin{bmatrix} F^i \\ \tau^i \end{bmatrix},$$

where $r^i = [r_x^i, r_y^i]^T$ is the inertial position of the i^{th} robot, θ^i the orientation, v^i the linear speed, ω_i the angular speed, τ^i the applied torque, F^i the applied force, m^i the mass, and J^i the moment of inertia for $i = 1, 2, \dots, N$. Applying output feedback linearisation about the hand position yields an approximate model for the position: $\ddot{x}^i = u^i$ (Lawton et al., 2003), which satisfies Assumption 2 as per Example 4. A formation pattern is a set

$$\mathcal{P} := \{x_*^1, \dots, x_*^N\},$$

where x_*^i denotes the desired constant location of the hand position of the i^{th} robot. Suppose the group of robots is required to transition through a sequence of formation patterns $\mathcal{P}_j$ which have been designed to avoid collisions. Lawton et al. (2003); Ren and Beard (2008) consider a bidirectional ring topology and propose the following control law which maintain the robots in the same shape as the destination pattern during the transition from one formation pattern to another:

$$\begin{aligned} u^i = & -K_g \tilde{x}^i - D_g \dot{\tilde{x}}^i \\ & -K_f(\tilde{x}^i - \tilde{x}^{i-1}) - D_f(\dot{\tilde{x}}^i - \dot{\tilde{x}}^{i-1}) \\ & -K_f(\tilde{x}^i - \tilde{x}^{i+1}) - D_f(\dot{\tilde{x}}^i - \dot{\tilde{x}}^{i+1}), \end{aligned} \quad (10)$$

where $\tilde{x}^i := x^i - x_*^i$, K_g and D_g are symmetric positive-definite matrices while K_f and D_f are symmetric positive-semidefinite matrices. Note that the indices are defined modulo N to observe the ring structure of the topology, whereby $x^{1-1} = x^N$ and $x^{N+1} = x^1$. It can be seen that the first two terms in (10) serve to drive the robot to reach its final position in the formation pattern. This accomplishes the reference forces design requirements of Theorem 7. The second two terms maintain the formation with the $i-1$ robot and the last two terms with the $i+1$ robot. These four terms guarantee that the formation pattern is an asymptotically stable equilibrium, i.e. Assumption 6 is satisfied. The coupled dynamics formation control law of (10) can be modified to take into account interrobot damping and actuator saturation constraints (Lawton et al., 2003; Ren and Beard, 2008).

Following the idea of the hardware experimental results of Lawton et al. (2003), the network manoeuvre from one formation pattern to another can be regarded as complete when each vehicle is within a pre-selected error distance tolerance ϵ of its destination; see (9). This error decreases with the increase of the time the fleet of vehicles is allocated to converge towards the destination's formation pattern. In the case of extreme source seeking, the formation pattern is updated recursively with the measurements of the field distribution by the extremum seeking controller.

Procedure 2 Suppose the transition to the formation pattern $\mathcal{P}_j$ has been completed, i.e. all the vehicles collected measurements of the field distribution Q at a distance no greater than $\epsilon := \frac{\nu}{L} > 0$ from their destinations; see Theorem 7. The measurements are exploited by the extremum seeking controller to estimate the gradient of Q around the local neighbourhood of where the vehicles are current situated. The next formation pattern $\mathcal{P}_{j+1}$ is then supplied by the controller to the vehicles. The reference and cooperative control law applied to each of the vehicle is adjusted accordingly to (10). The fleet of vehicles then proceeds to $\mathcal{P}_{j+1}$ in the prescribed formation.

5 Global extremum seeking

This section proposes a framework for performing global extreme source seeking of possibly non-convex and non-smooth field distributions. Deterministic iterative sampling-based algorithms are considered, which do not exploit information about gradients unlike those in the previous sections. A standing assumption is that the domain of the objective field distribution is a compact subset of $\mathbb{R}^n$, and this information is known to GPS-enabled mobile agents. The exploitation of parallelism to speed up the search for extrema is discussed at the end. Practical convergence to the extreme source is established in the face of bounded measurement noise. First, global extremum seeking from Khong et al. (2013) is reviewed.

5.1 Global extremum seeking algorithms

Consider the optimisation problem:

$$y^* := \max_{x \in \Omega} Q(x), \quad (11)$$

where $Q : \Omega \subset \mathbb{R}^n \rightarrow \mathbb{R}$ is a Lipschitz continuous field distribution which takes its *global* maximum value on $\mathcal{C} \subset \Omega$ and Ω is compact.

Suppose sampling is noisy as in Figure 1: $y_k = Q(x_k) + w_k$ with $|w_k| \leq \nu$ for some $\nu > 0$.

Assumption 8 Let $\delta \geq 0$ be a small number which characterises the accuracy of convergence. The discrete-time extremum seeking algorithm Σ satisfies the following: Given any $\mu > \delta$, there exists a $\nu > 0$ such that if $|w_k| \leq \nu$ for $k = 0, 1, \dots$, then for every $\epsilon > 0$ and $K \in \mathbb{N}$, there exists an $N \in \mathbb{N}$, $N > K$ for which $x_N \in \mathcal{C} + (\mu + \epsilon)\mathcal{B}$. In other words, a subsequence of the output of Σ , $\{x_k\}_{k=0}^\infty$ converges to a μ -neighbourhood of the set $\mathcal{C}$ of global maximisers of Q .

The reader is referred to (Khong et al., 2013, Section 3) for examples of recursive optimisation algorithms which satisfy the aforementioned assumption, such as the DIRECT method (Jones et al., 1993).

5.2 Multi-agent global extreme source seeking

Consider now the implementation of an optimisation algorithm satisfying Assumption 8 with N number of sensor-enabled mobile agents which satisfy Assumption 2. A main difference from the gradient-based work in Section 4 is that no specific formation is required to be maintained by the vehicles in this section and so each vehicle can be fed with a different reference velocity as illustrated in Figure 5.

Suppose the piecewise continuous inputs to the multiple vehicles are designed so that they collaboratively collect




Fig. 5. Global source seeking.

the required noise-corrupted samples of Q at the outputs x_k of Σ for $k = 0, 1, \dots$ as in Theorem 5; vehicle planning is investigated further in the forthcoming subsection.

Theorem 9 Consider the configuration of Figure 5 where the extremum seeker satisfies Assumption 8 and the vehicles satisfy Assumption 2. Given any $\mu > \delta$, where $\delta \geq 0$ is given in Assumption 8, there exists a $\nu > 0$ such that if $|w_k^i| \leq \nu$ for $i = 1, \dots, N$ and $k = 0, 1, \dots$, then a subsequence of the output of Σ , $\{x_k\}_{k=0}^\infty$, converges to $\mathcal{C} + \mu\mathcal{B}$.

PROOF. The result follows by the robustness property of the optimisation algorithm in Assumption 8. $\square$

Remark 10 Theorem 9 differs from the previous results in that the vehicles do not converge to the extreme source in a formation. Practically, the extremum seeking controller only runs for a finite amount of time before being terminated or switched off (a bound on the distance from the extremum can be calculated if the Lipschitz constant of Q is known; see Khong et al. (2013)). A point which lies close to the global maximum of a field may be one that has been found during a past iteration of the optimisation algorithm by one of the many vehicles when global extremum seeking of the sort in Theorem 9 is deployed. Should knowledge of the Lipschitz constant of Q be known, it can also be used to speed up the algorithm by taking out of consideration regions on which Q cannot possibly be optimal.

Remark 11 Note that there is a clear difference between the gradient optimisation based Theorem 7 and the global optimisation based Theorem 9. The former concludes asymptotic stability of local extrema of the field distribution. On the other hand, the latter states a subsequence attractivity type result for global extrema.

5.3 Vehicle assignment/routing

While Theorem 9 assumes that the measurements are collected by multiple vehicles, the challenging task of deciding which vehicle is to collect which point is left to the

designer, or perhaps a supervisory controller. This is the topic of this subsection. In reality, a vehicle takes time to travel from one point to another, and hence reducing the total distance travelled can increase the speed for localising a global extreme source. The communication between the vehicles and/or that with the supervisor is also considered.

5.3.1 Even partitioning

Suppose there is an abundance of vehicles and a maximum communication range/distance between two vehicles is $\Phi > 0$. A straightforward multi-agent scheme in this case is to divide the search space evenly into equal-size sub-regions such that the centre of any sub-region to those of its neighbours is no greater than Φ . In particular, these sub-regions are line segments in $\mathbb{R}$, rectangles in $\mathbb{R}^2$, and rectangular parallelepipeds or right cuboids in $\mathbb{R}^n$. See Figure 6 for a four-vehicle example in a bound-constrained $\Omega \subset \mathbb{R}^2$. It is assumed that there is no supervisory controller in this scenario and each of the vehicle is equipped with a global optimisation algorithm satisfying Assumption 8.




Fig. 6. Even districts. ‘V’ denotes vehicle.

Procedure 3 Each of the vehicles is to be positioned within its designated sub-region at the start of the search. As determined by the global extremum seeking algorithm, if a measurement to be collected falls within a sub-region, the corresponding vehicle is to travel towards it and perform the collection. Thereafter, the vehicle returns to the centre and information is exchanged between all the vehicles before the next iteration of the algorithm commences. Preference of vehicle can be assigned in advance should a sample point fall on the border between two or more sub-regions. Note that for some algorithms such as *DIRECT*, there may be more than one measurements to be collected within a sub-region, as in the bottom left area of Figure 6. The responsible vehicle can sample these points in a random order or following an optimised path returned by the traveling salesman problem or a special instance of the vehicle routing problem with one vehicle to be discussed below.

5.3.2 Vehicle routing problem

Suppose now that there is a central depot located at the centre of a bound-constrained $\Omega \subset \mathbb{R}^n$ in which N number of vehicles are available. Given K number of measurement destinations in Ω ($K > N$ possibly), these vehicles are to visit them, collect the required measurements, and return to the depot for information processing. The central depot serves as the supervisory controller which issues destination commands to each of the vehicle. By contrast to the previous partitioning scheme, the vehicles only need to store in their memory places to visit in a prescribed order and are not required to be endowed with the processing power to run the global extremum seeking algorithm.

The aforementioned setup thus gives rise to a vehicle routing problem, which is a combinatorial optimization and integer programming problem with the objective of servicing a number of customers with a fleet of vehicles (Toth and Vigo, 2002; Golden et al., 2008). See Figure 7 for a 3-vehicle example on $\Omega \subset \mathbb{R}^2$. Computing the optimal routes which minimise the total distance travelled by the agents is an NP-hard problem, with its decision version (given a set of routes, does there exist one that incurs a lower cost?) being NP-complete. Exact solutions can be computed using, for example, dynamic programming; see Toth and Vigo (2002); Laporte (1992); Golden et al. (2008). Various approximate and heuristic algorithms are also given in the references. Extensions to multiple depots can also be found. In this case, an underlying assumption is that these depots are scattered in the search space in such a way that communication is maintained between them.




Fig. 7. The vehicle routing problem

For a fully decentralised version of the problem, where each vehicle is equipped with a processing unit and the ability to communicate with every other agent in the search space regardless of the distance separation, after measurement collections the vehicles have no need to return to their starting positions at the commencement of the extremum seeking algorithm iteration. This problem can be tackled using the same exact or approximate algorithms for multi-depot vehicle routing, simply by relaxing the return-to-depot condition in the formulation.

6 Simulation examples

This section contains two examples on extreme source seeking with multiple agents. The first exploits the gradient-based extremum seeking framework of Section 4 while the second is based on the global extremum seeking control of Section 5.

6.1 Gradient-based extreme source seeking

Consider the following quadratic scalar field distribution

$$Q(z_1, z_2) := -(z_1 - 5)^2 + (z_2 - 7)^2 + 20,$$

which has a unique global maximum at $[5, 7]^T$. Suppose three mobile vehicles with single integrator dynamics are available, i.e. $\dot{x}^i = u^i$ for $i = 1, 2, 3$, and assume an undirected ring communication topology. The cooperative control and formation manoeuvre laws from Section 4.2 are adapted/simplified for these vehicles so that

$$u^i = 5\tilde{x}^i - (\tilde{x}^i - \tilde{x}^{i-1}) - (\tilde{x}^i - \tilde{x}^{i+1}),$$

where $\tilde{x} := x - x_*$ and x_* denotes the desired position of the i^{th} vehicle. The gradient descent method is employed

$$x_{k+1} = x_k + 0.2\nabla Q(x_k), \quad x_0 := [1, 1]^T, \quad (12)$$

together with the derivative estimator (6) with $h := 0.05$ as the extremum seeking controller of the form illustrated in Figure 2. Assume that the vehicles are initialised at the vertices of the right-angle triangle: $[1, 1]^T$, $[1+h, 1]^T$, and $[1, 1+h]^T$. The subsequent formation patterns are given by x_k , $x_k + [h, 0]^T$, and $x_k + [0, h]^T$ for $k = 1, 2, \dots$. A duration of 1s is allocated for the network of vehicles to manoeuvre from one formation pattern to another. The vehicles collect measurements of the field distribution at the end of the 1s and return this information to the extremum seeking controller as in Figure 4, which is used in the output update (12). The simulation results are as follow. Figure 8 shows the 2D positions of the vehicles as they converge to the global maximum point $[5, 7]^T$. The 2D components over time of the vehicles are plotted in Figure 9. These figures also demonstrate no collision occurred and the right-triangle formation is maintained during the formation manoeuvres.

6.2 Global extreme source seeking

Suppose the environment distribution is given by the Goldstein-Price test function:

$$\begin{aligned} Q(z_1, z_2) := & [1 + (z_1 + z_2 + 1)^2 \\ & (19 - 14z_1 + 3z_1^2 - 14z_2 + 6z_1z_2 + 3z_2^2)] \\ & [30 + (2z_1 - 3z_2)^2 \\ & (18 - 32z_1 + 12z_1^2 + 48z_2 - 36z_1z_2 + 27z_2^2)] \end{aligned}$$




Fig. 8. Vehicles' positions




Fig. 9. Vehicles' position components over time

with domain $-2 \leq z_1, z_2 \leq 2$. The function has multiple local minima and a global minimum 3 at $[0, -1]^T$. A surface plot of the Goldstein-Price function is provided in Figure 10.




Fig. 10. Goldstein-Price test function

Suppose 9 vehicles with single-integrator dynamics are deployed to localise the global minimum using the DIRECT optimisation based extremum seeking controller described in Section 5.1 and following the even partitioning protocol in Section 5.3.1. Further assume that the measurements are corrupted by a uniformly distributed random variable with support $[-1, 1]$. Suppose that each iteration of DIRECT takes 1 unit of time to be completed, which includes the time for vehicles' movements to places where measurements are to be retrieved, information exchange, and data processing. After 9 iterations and 83 sample collections, an estimate of the global minimum value made by the DIRECT method is 3.0325,

which was collected at $[0, -1.037]^T$. The estimates of the extremum seeking controller over time are plotted in Figure 11.




Fig. 11. The minimum value located by the extremum seeking controller over time

7 Conclusions

This paper proposes two unified approaches to the extreme source seeking problem with multiple agents. First, gradient-based local optimisation is considered in which a network of vehicles gradually converges to the source of a field environment. The formation is cooperatively maintained for gradient estimation with bounded errors. Global extremum seeking in the face of noise-corrupted measurements is also examined. Examples of optimisation algorithms and cooperative control algorithms are provided and the vehicle distribution/routing problem is discussed.

References

- Bai, H., Arcak, M., and Wen, J. (2011). *Cooperative Control Design: A Systematic, Passivity-Based Approach*. Springer.
- Biyik, E. and Arcak, M. (2008). Gradient climbing in formation via extremum seeking and passivity-based coordination rules. *Asian Journal of Control*, 10(2):201–211.
- Cochran, J. and Krstic, M. (2009). Nonholonomic source seeking with turning of angular velocity. *IEEE Transactions on Automatic Control*, 54(4):717–731.
- Fiorelli, E., Leonard, N. E., Bhatta, P., Paley, D. A., Bachmayer, R., and Fratantoni, D. M. (2006). Multi-AUV control and adaptive sampling in Monterey Bay. *IEEE Journal of Oceanic Engineering*, 31(4):935–948.
- Golden, B. L., Raghavan, S., and Wasil, E. A., editors (2008). *The Vehicle Routing Problem: Latest Advances and New Challenges*. Springer.
- Jones, D. R., Pertunnen, C. D., and Stuckman, B. E. (1993). Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Applications*, 79(1):157–181.
- Khalil, H. K. (2002). *Nonlinear Systems*. Prentice Hall, 3rd edition.
- Khong, S. Z., Nešić, D., Tan, Y., and Manzie, C. (2013). Unified frameworks for sampled-data extremum seeking control: global optimisation and multi-unit systems. *Automatica*, 49(9):2720–2733.
- Laporte, G. (1992). The vehicle routing problem: An overview of exact and approximate algorithms. *European Journal of Operational Research*, 59:345–358.
- Lawton, J. R. T., Beard, R. W., and Young, B. J. (2003). A decentralized approach to formation maneuvers. *IEEE Transactions on Robotics and Automation*, 19(6):933–941.
- Leonard, N. E. and Fiorelli, E. (2001). Virtual leaders, artificial potentials and coordinated control of groups. In *Conference on Decision and Control*, pages 2968–2973.
- Leonard, N. E., Paley, D. A., Davis, R. E., Fratantoni, D. M., Lekien, F., and Zhang, F. (2010). Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in Monterey Bay. *Journal of Field Robotics*, 27(6):718–740.
- Mayhew, C. G., Sanfelice, R. G., and Teel, A. R. (2007). Robust source-seeking hybrid controllers for autonomous vehicles. In *American Control Conference*.
- Mayhew, C. G., Sanfelice, R. G., and Teel, A. R. (2008). Robust source-seeking hybrid controllers for nonholonomic vehicles. In *American Control Conference*.
- Nešić, D., Teel, A. R., and Kokotović, P. V. (1999). Sufficient conditions for stabilization of sampled-data nonlinear systems via discrete-time approximations. *Systems & Control Letters*, 38:259–270.
- Ögren, 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.
- Ren, W. and Beard, R. W. (2008). *Distributed Consensus in Multi-vehicle Cooperative Control*. Springer.
- Reyhanoglu, M., van der Schaft, A., McClamroch, N. H., and Kolmanovsky, I. (1999). Dynamics and control of a class of underactuated mechanical systems. *IEEE Transactions on Automatic Control*, 44(9):1663–1671.
- Teel, A. R. (2000). Lyapunov methods in nonsmooth optimization, part I: Quasi-Newton algorithms for Lipschitz, regular functions and part II: persistently exciting finite differences. In *Proceedings of IEEE Conference on Decision and Control*.
- Teel, A. R. and Popović, D. (2001). Solving smooth and nonsmooth multivariable extremum seeking problems by the methods of nonlinear programming. In *Proceedings of American Control Conference*, volume 3, pages 2394–2399.
- Toth, P. and Vigo, D. (2002). *The Vehicle Routing Problem*. SIAM.
- Zhang, C., Siranosian, A., and Krstić, M. (2007). Extremum seeking for moderately unstable systems and for autonomous vehicle target tracking without position measurements. *Automatica*, 43:1832–1839.