Using Radial Basis Functions to Approximate the LQG-Optimal Event-Based Sampling Policy

Thelander Andrén, Marcus

Published in:
2019 18th European Control Conference (ECC)

DOI:
10.23919/ECC.2019.8795838

2019

Document Version:
Peer reviewed version (aka post-print)

Using Radial Basis Functions to Approximate the LQG-Optimal Event-Based Sampling Policy

Marcus Thelander Andrén

***Abstract*— A numerical method using radial basis functions (RBF) has been developed to find the optimal event-based sampling policy in an LQG problem setting. The optimal sampling problem can be posed as a stationary partial differential equation with a free boundary, which is solved by reformulating the optimal RBF approximation as a linear complementarity problem (LCP). The LCP can be efficiently solved using any quadratic program solver, and we give guarantees of existence and uniqueness of the solution. The RBF method is validated numerically, and we showcase what the different types of optimal policies look like for 2D systems.**

***Index Terms*— Event-based sampling, LQG-optimal control, sampled-data control, radial basis functions**

I. INTRODUCTION

In the field of sampled-data systems, the concept of *event-based control* is to trigger sampling and actuation based on the behavior of controlled variables, in contrast to traditional triggering based on a periodic timer. The motivation is a more resource-efficient design, with potential to save e.g. energy, network bandwidth and computations. This potential was clearly demonstrated in the early works [1] and [2], and it has since become a very active field of research [3]–[6].

There are two degrees of freedom to consider for event-based controllers; (i) the intersample behaviour and (ii) the sampling policy. For a closed-loop sampled-data system of the form in Fig. 1, (i) corresponds to the design of the sampler $\mathcal{S}$, the hold circuit $\mathcal{H}$ and the discrete-time controller $\bar{K}$ while (ii) is to design a policy which decides the sequence of sampling times $\{t_i\}_{i\in\mathbb{N}_0}$. Finding an optimal co-design for this system is generally considered a difficult task, and many previous works have therefore focused on sub-optimal policies [7]–[10] or data-driven approaches [11], [12] shown to outperform their periodic counterparts.

Recently however, an $\mathcal{H}_2$-optimal design of $\mathcal{S}$,$\bar{K}$ and $\mathcal{H}$ was presented in [13], which remains optimal for any given, uniformly bounded, sequence of sampling times. Furthermore, the optimal structure was shown to behave equivalently to the system shown in Fig. 2. This remarkable result was later applied to event-based control in [14] and [15], wherein the latter it was formally proven that the structure remains optimal in an LQG setting with sampling times depending on the controlled variables. The implication is that the co-design problem is separable in the LQG case, and that the remaining problem is to find the sampling policy which optimizes the trade-off between LQG cost and sampling rate.

This work has been supported by the Swedish Research Council. The author is with the Department of Automatic Control, Lund University, Sweden, and is a member of the LCCC Linnaeus Center and the ELLIIT Excellence Center at Lund University. E-mail: marcus@control.lth.se

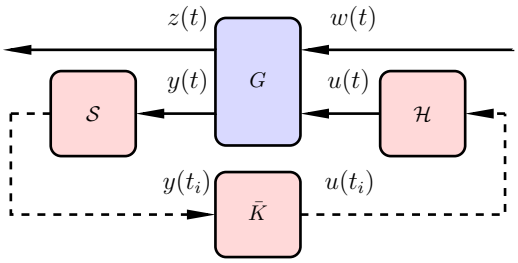

Fig. 1. An LTI system G in feedback with a sampled-data controller consisting of a sampler $\mathcal{S}$, a hold circuit $\mathcal{H}$ and a discrete-time controller $\bar{K}$. Solid lines represent continuous-time signals, while dashed lines represent discrete-time signals.

In our previous work [16] we used the framework for optimal impulse control in [17, Paper I & II] to formulate the optimal event-based sampling policy as the solution to a stationary partial-differential equation (PDE) with free boundary. Solving this type of PDE is a challenging problem, which generally requires numerical methods. To handle the free boundary, we proposed in [16] a finite-difference method based on simulating a time-dependent version of the PDE from some initial guess, and then extract the solution once the simulation reached stationarity. However, the introduction of time-dependence brought several issues, such as errors due to time discretization, dependency on initial guess and ambiguous conditions for when the simulation is sufficiently close to stationarity. Our contribution in this paper is to derive a more efficient numerical method based on radial basis functions (RBF) [18], which solves the stationary problem directly. The method is inspired by solutions to similar free boundary problems in mathematical finance [19], where the optimization of RBF weights can be formulated as a *linear complementarity problem* [20]. The optimal RBF weights are then easily obtained by solving a quadratic program (QP), and we give guarantees for existence and uniqueness of the optimal solution. The method is numerically validated against an analytic solution of the PDE for a special case, and we use the method to characterize different possible types of optimal sampling policies in the 2D case. The results could be used to guide future designs of near-optimal but simpler event-based sampling policies. To highlight this, we also showcase a numerical example where we compare the performance between the optimal sampling policy and a much simpler heuristic policy.

II. PROBLEM FORMULATION

A. Setup and Goal

We consider the problem of finding the optimal event-based sampling policy for the closed-loop system shown in

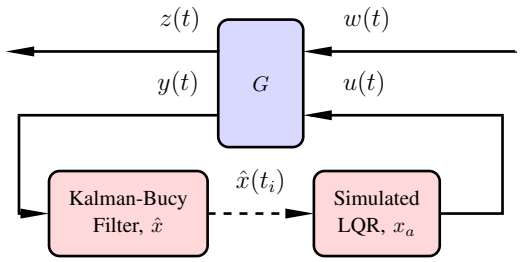

Fig. 2. The closed-loop system for which we are designing an optimal sampling policy. It consists of an LTI system G and a Kalman-Bucy filter on the sensor side that intermittently transmits its estimate $\hat{x}$ to an LQR controller simulating the closed-loop system on the actuator side.

Fig. 2. It consists of a linear time-invariant (LTI) continuous-time plant G, a Kalman-Bucy filter operating on the measurement y on the sensor side and an LQR controller on the actuator side which generates the control signal u based on a simulation of G. The plant is subject to disturbances in the form of a vector Gaussian white process w with unit intensity. The controlled output z is a linear combination of the plant state and control signal, and expresses the closed-loop performance in terms of the infinite-horizon LQG cost

$$J_z \triangleq \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}[\int_0^T z(t)^\intercal z(t)dt]. \tag{1}$$

At sampling times $\{t_i\}_{i\in\mathbb{N}_0}$, the Kalman-Bucy filter transmits its estimate $\hat{x}$ of the plant state vector to the actuator side, where the simulated plant state vector x_a is reset to $\hat{x}$, i.e. $x_a(t_i) = \hat{x}(t_i)$. The motivation behind this structure is that it is an equivalent representation of the optimal controller structure for the closed loop system in Fig. 1. For details on this connection, we refer to the original derivation in [13] and the subsequent works [15] and [16].

The goal is to design an event-based sampling policy which achieves an efficient trade-off between the transmission of $\hat{x}$ (incurring costs in e.g. energy and network bandwidth) and the closed-loop performance J_z. To this end we define the average sampling rate as

$$f \triangleq \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}[\sum_{i=0}^{\infty} 1_{t_i\leq T}],$$

where the sum counts the number of sampling events up to time T. Adding a fixed penalty $\rho \geq 0$ per sample, our goal is to find a sampling policy which minimizes the total cost

$$J_z + \rho f. \tag{2}$$

B. Models

The plant G has the following n-dimensional realization

$$G: \begin{cases} \dot{x}(t) = Ax(t) + B_w w(t) + B_u u(t), \\ z(t) = C_z x(t) + D_{zu} u(t), \\ y(t) = C_y x(t) + D_{yw} w(t), \end{cases} \tag{3}$$

which is assumed to satisfy the standard conditions on well-posedness of $\mathcal{H}_2$ control [21, Sec. 14.5]. For a given plant (3), we can compute the corresponding Kalman-Bucy gain L and LQR gain F by solving the two algebraic Riccati equations

$$\begin{cases} A^\intercal X + XA + C_z^\intercal C_z - F^\intercal(D_{zu}^\intercal D_{zu})F = 0, \\ F = -(D_{zu}^\intercal D_{zu})^{-1}(B_u^\intercal X + D_{zu}^\intercal C_z), \end{cases}$$

$$\begin{cases} AY + YA^\intercal + B_w B_w^\intercal - L(D_{yw}D_{yw}^\intercal)L^\intercal = 0, \\ L = -(YC_y^\intercal + B_w D_{yw}^\intercal)(D_{yw}D_{yw}^\intercal)^{-1}. \end{cases}$$

The Kalman-Bucy filter in Fig. 2 is then given by

$$\dot{\hat{x}}(t) = A\hat{x}(t) + B_u u(t) - L(y(t) - C_y\hat{x}(t)). \tag{4}$$

The LQR controller on the actuator side, which is based on an intermittently reset simulation of the plant, is given by

$$\begin{cases} \dot{x}_a(t) &= (A + B_u F)x_a(t), \qquad x_a(t_i) = \hat{x}(t_i), \\ u(t) &= Fx_a(t). \end{cases} \tag{5}$$

C. The Optimal Sampling Problem

In the degenerate case $\rho = 0$ there is no cost on sampling, and the optimal sampling policy thus becomes trivial; sample infinitely fast (i.e. $x_a(t) = \hat{x}(t)$, $\forall t$) and retain the continuous-time LQG controller. We will then achieve the minimum cost $\gamma_0 \triangleq \min J_z$, given by [21, Thm. 14.7]

$$\gamma_0 = \mathrm{Tr}(B_w^\intercal X B_w) + \mathrm{Tr}(C_z Y C_z^\intercal) + 2\mathrm{Tr}(XAY). \tag{6}$$

The cost γ_0 is the fundamental lower bound on J_z, and no policy can achieve a better closed-loop performance.

When $\rho > 0$ it is clear that any sampling policy minimizing (2) must have a finite average sampling rate $f > 0$ (the closed-loop system can be unstable for $f = 0$). The performance of the closed-loop system is then fundamentally dependent on the inter-sampling error $\tilde{x} \triangleq \hat{x} - x_a$ between the Kalman-Bucy estimate $\hat{x}$ and the state x_a of the LQR simulation. From (4) and (5) we compute the dynamics

$$\dot{\tilde{x}} = A\tilde{x}(t) + v(t), \quad \tilde{x}(t_i) = 0, \tag{7}$$

where the innovation $v = -L(y - C_y\hat{x})$ of the Kalman-Bucy filter is a vector Gaussian white process with intensity $LD_{yw}(LD_{yw})^\intercal \triangleq R \succ 0$. Note in (7) that the action of sampling corresponds to resetting the error $\tilde{x}$ to zero. The fundamental role of $\tilde{x}$ becomes apparent from the results in [15, Thm. 1], where it is shown that

$$J_z = \gamma_0 + \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}[\int_0^T \tilde{x}(t)^\intercal Q\tilde{x}(t)dt], \tag{8}$$

with $Q = (D_{zu}F)^\intercal D_{zu}F$. We conclude from (8) that the cost J_z is kept small if we choose to sample often, and thus reset $\tilde{x}$ to zero. However, we also incur the fixed cost ρ each time we sample. This is the trade-off our sampling policy should optimize. From (8) we note that the constant γ_0 can be disregarded when minimizing (2), and we are now ready to formally define the optimal sampling problem:

OSP$_\rho$: Consider the closed-loop system in Fig. 2, with a Kalman-Bucy filter governed by (4) and an LQR controller based on a simulated plant according to (5). Denote

$$J_{\tilde{x}} \triangleq \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}[\int_0^T \tilde{x}(t)^\intercal Q\tilde{x}(t)dt]. \tag{9}$$

The optimal sampling problem is then to design a sampling policy which minimizes the cost

$$J \triangleq J_{\tilde{x}} + \rho f. \tag{10}$$

Note that **OSP**$_\rho$ can be seen as an optimal stopping problem, where we search for a threshold on $\tilde{x}$ which optimizes the trade-off between expected average quadratic cost and the deterministic cost ρ for resetting, see Fig. 3.

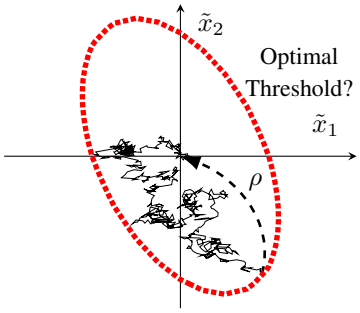

Fig. 3. The optimal sampling problem (here in 2D) can be seen as searching for a threshold (dashed) on the error $\tilde{x}$ (solid), from where we should reset $\tilde{x}$ to the origin and incur the fixed cost ρ.

Designing a sampling policy based on $\tilde{x}$ is possible by using a copy of (5) on the sensor side.

Next we will characterize the optimal policy for $\mathbf{OSP}_\rho$ as the solution of a free boundary PDE, which is the focus of the proposed RBF method of this paper.

D. PDE-Formulation of the Optimal Sampling Problem

In [17, Paper I & II] it was shown that $\mathbf{OSP}_\rho$ can be posed as the solution to a stationary partial differential equation (PDE) with a free boundary. Here we outline how this PDE is derived, and refer to [17, Paper II] for a detailed proof.

Let j denote the accumulated cost up until time T, i.e.

$$j \triangleq \mathbb{E}[\int_0^T \tilde{x}^\intercal(t)Q\tilde{x}(t)dt] + \rho\mathbb{E}[\sum_{i=0}^{\infty} 1_{t_i \leq T}].$$

In the spirit of dynamic programming, we are searching for a bounded *relative value function* $V : \tilde{x} \to \mathbb{R}$ and a scalar $\underline{J}$ satisfying

$$\underline{J}T \leq j + \mathbb{E}[V(\tilde{x}(T)) - V(\tilde{x}(0))], \quad \forall T \in [0, \infty). \quad (11)$$

If such a pair is found, then division by T and taking the limit $T \to \infty$ gives the lower bound $\underline{J} \leq J$ on the cost in (10). Note that if (11) is an equality for all T, then $\underline{J} = J$.

We now continue by deriving conditions on V and $\underline{J}$ such that (11) is satisfied. Denote $\delta V \triangleq V(\tilde{x}(T)) - V(\tilde{x}(0))$ and let $I(T)$ be the set of indices of the sampling times occurring on $[0, T]$. Furthermore, let $\mathcal{T} = [0, T] \setminus \{t_i\}_{i \in I(T)}$ denote the intervals of time between samples. By using (7) and Itō's formula, we can compute the expected change in V as

$$\begin{aligned}\mathbb{E}[\delta V] &= \mathbb{E}[\int_{\mathcal{T}} dV] + \mathbb{E}[\sum_{i \in I(T)} V(0) - V(\tilde{x}(t_i^-))] \\ &= \mathbb{E}[\int_{t \in \mathcal{T}} (\tilde{x}^\intercal A^\intercal \nabla V + \frac{1}{2}\mathrm{Tr}(R\nabla^2 V))dt] \\ &+ \mathbb{E}[\sum_{i \in I(T)} V(0) - V(\tilde{x}(t_i^-))],\end{aligned}$$

where $\tilde{x}(t_i^-)$ denotes $\lim_{t \uparrow t_i} \tilde{x}(t)$. If V and $\underline{J}$ satisfy

$$\tilde{x}^\intercal Q\tilde{x} + \tilde{x}^\intercal A^\intercal \nabla V + \frac{1}{2}\mathrm{Tr}(R\nabla^2 V) \geq \underline{J}, \quad \forall \tilde{x}, \quad (12)$$

$$\rho + V(0) - V(\tilde{x}) \geq 0, \quad \forall \tilde{x}, \quad (13)$$

then it follows that

$$\begin{aligned}j + \mathbb{E}[\delta V] &= \mathbb{E}[\int_{t \in \mathcal{T}} (\tilde{x}^\intercal Q\tilde{x} + \tilde{x}^\intercal A^\intercal \nabla V + \frac{1}{2}\mathrm{Tr}(R\nabla^2 V))dt] \\ &+ \mathbb{E}[\sum_{i \in I(T)} \rho + V(0) - V(\tilde{x}(t_i^-))] \geq \underline{J}\mathbb{E}[\int_{t \in \mathcal{T}} dt] = \underline{J}T,\end{aligned}$$

i.e. (12) and (13) implies the inequality (11). If we put stricter conditions on V, enforcing that (12) and (13) are equalities for those $\tilde{x}$ where sampling is not triggered and triggered respectively, then we arrive at the following result:

Theorem 1 ([17, Paper II, Thm. 1]): Suppose a bounded, $\mathcal{C}^2$, function V and a constant $\underline{J} = J$ satisfy (12) and (13), with equality in at least one of them for all $\tilde{x}$. Then the optimal cost in $\mathbf{OSP}_\rho$ is J, and an optimal sampling policy is to trigger sampling whenever equality holds in (13), i.e. the sequence of optimal sampling times $\{t_i\}$ are given by $t_i = \min\{t > t_{i-1} : \rho + V(0) - V(\tilde{x}(t)) = 0\}$.

Proof: See proof of Theorem 1 in [17, Paper II]. ■

The conditions on V and J in Theorem 1 can also be more compactly formulated as

$$\begin{aligned}\min\{\tilde{x}^\intercal Q\tilde{x} - J + \tilde{x}^\intercal A^\intercal \nabla V + \frac{1}{2}\mathrm{Tr}(R\nabla^2 V), \\ \rho + V(0) - V(\tilde{x})\} = 0, \quad \forall \tilde{x}.\end{aligned} \quad (14)$$

For a given J, this is a stationary PDE with a free boundary. The free boundary is implicitly given by those $\tilde{x}$ where the expression minimizing (14) changes. This class of PDE's typically arises in optimal stopping problems, e.g. in pricing of American options in finance [19].

While an analytic solution is available for the special case $A = 0$ (see [17] and [16]), there is little hope of finding such a solution for the general case with $A \neq 0$. Instead we turn to numerical methods to approximate the solution to some required precision. In the next section we derive an RBF method which solves the stationary problem (14) directly for some given value of J.

III. A Radial Basis Function Approximation

A. Preliminaries

We start with some simple observations in (14). First, linear transformations can always be performed on $\tilde{x}$ such that $R = I$ and Q is a diagonal matrix in the transformed variable. Second, the choice of $V(0)$ is non-consequential since it is only a reference value, and can thus be set to $V(0) = -\rho$ to eliminate the explicit dependence of ρ. With these steps performed, we write (14) as

$$\begin{aligned}-V(\tilde{x}^\intercal Q\tilde{x} - J + \tilde{x}^\intercal A^\intercal \nabla V + \frac{1}{2}\Delta V) = 0, \quad \forall \tilde{x}, \\ -V \geq 0, \;\; \tilde{x}^\intercal Q\tilde{x} - J + \tilde{x}^\intercal A^\intercal \nabla V + \frac{1}{2}\Delta V \geq 0, \;\; \forall \tilde{x}.\end{aligned} \quad (15)$$

Henceforth the PDE of the form in (15) will be used.

B. Approximation using RBFs

Our aim is to approximate V as a weighted sum of radial basis functions $\phi_j(\tilde{x}) : \tilde{x} \to \mathbb{R}$, where each basis function is radially symmetric and centered at one of a set of N given *collocation points* $\{\tilde{x}_j\}_{j=1}^N$. The approximation $\hat{V}$ is

$$V(\tilde{x}) \approx \hat{V}(\tilde{x}) \triangleq \sum_{j=1}^{N} \alpha_j \phi_j(\tilde{x}), \quad (16)$$

where $\{\alpha_j\}_{j=1}^N$ is a set of weights to be determined. The RBF approximation is *mesh free*, meaning that we can choose the set of collocation points freely in the state space, not constrained to a uniform grid as for example finite-difference approximations.

While there are many choices for basis functions in the literature [18], a popular choice, which we will use here, are Gaussian basis functions

$$\phi_j(\tilde{x}) \triangleq \exp(-c||\tilde{x} - \tilde{x}_j||_2^2).$$

The parameter $c > 0$ is known as a *shape parameter*, and determines the decay rate of the basis functions. It is typically chosen as a trade-off between accuracy and numerical stability, where a small value of c often will improve the accuracy at the price of ill-conditioning [18].

With the choice of Gaussian basis functions, we can analytically compute the gradient and Laplacian of $\hat{V}$ as

$$\nabla\hat{V} = \sum_{j=1}^{N} \alpha_j \nabla\phi_j(\tilde{x}) = -2c\sum_{j=1}^{N} \alpha_j(\tilde{x} - \tilde{x}_j)\phi_j(\tilde{x}), \tag{17}$$

$$\Delta\hat{V} = \sum_{j=1}^{N} \alpha_j \Delta\phi_j(\tilde{x}) = 2c\sum_{j=1}^{N} \alpha_j(2c\|\tilde{x} - \tilde{x}_j\|_2^2 - n)\phi_j(\tilde{x}). \tag{18}$$

Thus, inserting $\hat{V}$ into (15) yields

$$\begin{aligned} -(\sum_{j=1}^{N} \alpha_j\phi_j(\tilde{x}))(\tilde{x}^\intercal Q\tilde{x} - J + \sum_{j=1}^{N} \alpha_j\lambda_j(\tilde{x})) = 0, \\ -\sum_{j=1}^{N} \alpha_j\phi_j(\tilde{x}) \geq 0, \quad \tilde{x}^\intercal Q\tilde{x} - J + \sum_{j=1}^{N} \alpha_j\lambda_j(\tilde{x}) \geq 0, \end{aligned} \tag{19}$$

where $\lambda_j(\tilde{x})$ is given by

$$\lambda_j(\tilde{x}) \triangleq c(2c\|\tilde{x} - \tilde{x}_j\|_2^2 - 2\tilde{x}^\intercal A^\intercal(\tilde{x} - \tilde{x}_j) - n)\phi_j(\tilde{x}).$$

Note that if (19) would be satisfied for all $\tilde{x}$, then $\hat{V}$ would in fact be a true solution to (15). Since we generally can not guarantee this, we instead relax the condition and specify that (19) must be satisfied at the collocation points $\{\tilde{x}_j\}_{j=1}^N$. We then get a system of N equations subject to $2N$ inequalities, and the goal is now to find a set of weights $\{\alpha_j\}_{j=1}^N$ such that they are all satisfied. In the next section we will derive an equivalent QP formulation of this problem.

IV. Computing the RBF Weights

A. The Linear Complementarity Problem

Inspired by how American option prices are approximated in [19], we proceed by formulating a linear complementarity problem (LCP) which ensures that (19) is satisfied at all collocation points. To this end we define the vectors

$$\alpha \triangleq \left[\alpha_1, ..., \alpha_j, ..., \alpha_N\right]^\intercal,$$

$$\beta \triangleq \left[\tilde{x}_1 Q\tilde{x}_1 - J, ..., \tilde{x}_j Q\tilde{x}_j - J, ..., \tilde{x}_N Q\tilde{x}_N - J\right]^\intercal,$$

and the matrices $\Phi, \Lambda \in \mathbb{R}^{N\times N}$, whose elements on the ith row and jth column are given by

$$\Phi_{i,j} \triangleq \phi_j(\tilde{x}_i), \quad \Lambda_{i,j} \triangleq \lambda_j(\tilde{x}_i). \tag{20}$$

The condition that (19) should hold for all collocation points can be expressed as

$$(-\Phi\alpha)_i(\Lambda\alpha + \beta)_i = 0, \quad \forall i = 1...N, \tag{21}$$

$$\text{s.t } -\Phi\alpha \geq 0, \quad \Lambda\alpha + \beta \geq 0, \tag{22}$$

where the inequalities are considered element-wise, and $(\cdot)_i$ denotes the ith element of a vector. Since the factors in (21) are non-negative, this is equivalent to

$$\begin{aligned} (-\Phi\alpha)^\intercal(\Lambda\alpha + \beta) = 0, \\ \text{s.t } -\Phi\alpha \geq 0, \quad \Lambda\alpha + \beta \geq 0. \end{aligned}$$

Finally, let $z \triangleq -\Phi\alpha$, i.e $\alpha = -\Phi^{-1}z$ (Gaussian basis functions guarantees $\Phi \succ 0$ [18]), which gives

$$\begin{aligned} z^\intercal(Mz + \beta) = 0, \\ \text{st.} \quad z \geq 0,\ Mz + \beta \geq 0, \end{aligned} \tag{23}$$

where $M = -\Lambda\Phi^{-1}$. Finding a z satisfying (23) is an LCP, and is equivalent to the QP

$$\begin{aligned} \min_z\ z^\intercal(Mz + \beta), \\ \text{st.} \quad z \geq 0,\ Mz + \beta \geq 0, \end{aligned} \tag{24}$$

with the minimum objective zero. The problem (24) is efficiently solved using any QP solver, and after obtaining a solution z^* we simply compute the weights as $\alpha = -\Phi^{-1}z^*$.

B. Uniqueness of Solution

Guarantees of existence and uniqueness of the solution to (24) are given by the following theorem:

Theorem 2: For the QP (24) with $M = -\Lambda\Phi^{-1}$, there exists a finite $\underline{c} > 0$ such that for any shape parameter $c > \underline{c}$, (24) has a unique solution.

Proof: The proof is given in the appendix. ■

While Theorem 2 states the existence of a finite lower bound on c such that (24) has a unique solution, it does not provide this bound explicitly. The result of Theorem 2 is based on the fact that an LCP of the form (23) has a unique solution for every β iff M is a P-matrix (a matrix with positive principal minors) [20]. Thus in order to verify that the value of c is chosen sufficiently large for our problem, we need to verify if M is a P-matrix or not. Unfortunately, this task is an NP-hard problem, and a straightforward evaluation of the $2^N - 1$ principal minors typically requires some $N^3 2^N$ operations [22]. This means we can only practically perform this verification for small N. However, there exist several sufficient conditions of less computational complexity that can be used. For instance, M is a P-matrix if $M + M^\intercal \succ 0$, which requires some N^3 operations to verify. For more sufficient conditions, see [23] and references therein.

C. Summary of Method

Here follows a summary of our proposed method to compute an approximate solution to **OSP**$_\rho$[1]:

1) Given the plant model (3), extract A and compute the innovation intensity R and weight matrix Q.
2) Perform a linear transformation of the error state $\tilde{x}$ such that $R = I$ in the transformed state.
3) Pick a cost J.
4) Pick a set of collocation points $\{\tilde{x}_j\}_{j=1}^N$ and a shape parameter c for the RBF approximation $\hat{V}$.
5) Compute $M = -\Lambda\Phi^{-1}$ and verify according to Section IV-B that (24) has a unique solution. If not, increase c.
6) Solve (24) using any QP solver, and compute the RBF weights as $\alpha = -\Phi^{-1}z^*$.
7) The approximation of the optimal sampling policy is now to sample whenever $\hat{V}(\tilde{x}) = 0$. The cost per sampling action ρ is given by $\hat{V}(0)$.

Remark 1: Depending on the size of the interpolation errors in the RBF approximation $\hat{V}$, the threshold $\hat{V}(\tilde{x}) = 0$

[1]Demo code is available at https://gitlab.control.lth.se/marcus/rbf-approx

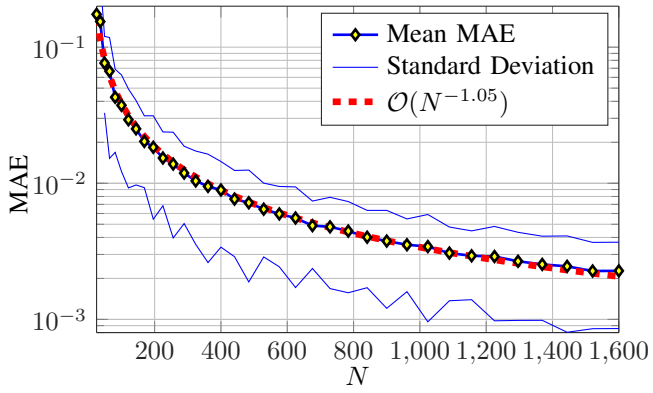

Fig. 4. Mean and standard deviation of MAE of RBF approximation over 50 randomized **OSP**$_\rho$ for different number of collocation points. A curve of order $\mathcal{O}(N^{-1.05})$ is plotted for reference.

might not be a single coherent surface as we would expect from the true solution. In that case, it is instead preferable to consider the threshold $\hat{V}(\tilde{x}) = \epsilon$, where ϵ is a small negative value close to zero chosen such that the threshold is coherent.

V. Numerical Evaluation

In this section we consider numerical validation of the proposed RBF method, and showcase the optimal policies for different classes of systems. We also present an example where we compare the performance of periodic sampling, a simple heuristic event-based sampling policy and the optimal event-based sampling policy.

A. Validation

The proposed RBF method is numerically validated using the analytic solution for the special case when $A = 0$. The solution was derived in [17, Paper II], and is given by

$$V(\tilde{x}) = -\frac{1}{4}(\max\{0, 2\sqrt{\rho} - \tilde{x}^\intercal P \tilde{x}\})^2, \tag{25}$$

where P is the unique solution to the Riccati-like equation

$$PRP + \frac{1}{2}\mathrm{Tr}(RP)P = Q. \tag{26}$$

In the validation, we randomize 50 versions of **OSP**$_\rho$ with a 2D plant, and for each version compute the maximum absolute error (MAE) $\|V - \hat{V}\|_\infty$ of the RBF approximations using $N = 5^2, 6^2, ..., 40^2$ uniformly distributed collocation points. Since we can assume $R = I$ and $Q = \mathrm{diag}(q_1, q_2)$ without loss of generality (see Section III-A), the **OSP**$_\rho$ is uniquely described by the parameters q_1, q_2 and ρ. These parameters were chosen randomly for the 50 **OSP**$_\rho$, with $\rho \in [0.01,\ 1]$ and $q_1, q_2 \in [1,\ 10]$.

The mean and standard deviation of the MAE over all 50 versions are shown in Fig. 4. Additionally, the RBF approximation for one example **OSP**$_\rho$ from the validation is shown in Fig. 5. The results indicate that the RBF approximation converges to the true solution, with an order of accuracy of roughly $\mathcal{O}(N^{-1})$.

B. Types of Policies for 2D Systems

While the optimal policies are known to be elliptic in the special case $A = 0$, it has been largely unknown what these policies look like in the general case. In [16], we observed the optimal policies for a couple of choices of A, and concluded that they are not necessarily convex. Here we make a more

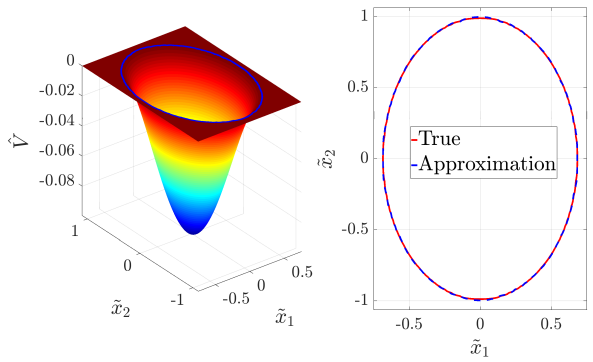

Fig. 5. Validation example using $A = 0$, $q_1 = 3$, $q_2 = 1$ and $\rho = 0.1$. The RBF approximation $\hat{V}$ (left) has $50^2 = 2,500$ collocation points and an MAE of $4.6 \cdot 10^{-4}$. The approximation of the optimal policy compares well to the true optimal policy (right).

thorough characterization of the types of optimal policies that are possible for different systems in the 2D case.

Based on the different types of possible phase portraits of the expected trajectory, we have chosen to investigate systems with a saddle point, a double integrator, a center, a star, a node and a spiral. Since we are interested in the impact of A, we keep $R = Q = I$ and $J = 1$ fixed throughout. The A-matrices for the different systems are given by

$$\text{Saddle: } A = \begin{bmatrix} 0 & 15 \\ 15 & 0 \end{bmatrix}, \quad \text{Integrator: } A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix},$$

$$\text{Center: } A = \begin{bmatrix} 0 & 15 \\ -15 & 0 \end{bmatrix}, \quad \text{Star: } A = \begin{bmatrix} 15 & 0 \\ 0 & 15 \end{bmatrix},$$

$$\text{Node: } A = \begin{bmatrix} 15 & 10 \\ 10 & 15 \end{bmatrix}, \quad \text{Spiral: } A = \frac{15}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}.$$

The resulting optimal polices are shown in Fig. 6. Perhaps the most striking observation is that most of the policies either are, or are well-approximated by, circles. The fact that the policies for the center, star and spiral are all circles becomes apparent when noting that their phase portraits are symmetric. Less obvious are the shapes of the double integrator and node, which despite their asymmetric phase portraits have optimal policies which could be well approximated by circles. This suggests that simple heuristic sampling policies parametrized by an ellipse could be designed such that near-optimal performance is achieved in these cases.

The only notably different case is the system with a saddle point, which results in a non-convex policy. This curious result also appears in higher dimensions, as shown in the 3D case in Fig. 7. While the non-convex policy itself is radically different from the other cases, it remains to be quantified how much better this policy actually can perform over, say, an elliptic policy. This is investigated in the next section.

C. Performance Comparison

Here we compare the trade-off between the cost $J_{\tilde{x}}$ and the average sampling period $h_{\text{avg}} \triangleq 1/f$ for three different sampling policies:

(a) Periodic sampling.
(b) The optimal sampling policy, approximated using the proposed RBF method.
(c) A simple heuristic policy, where sampling is triggered whenever $\|\tilde{x}\|_2 \geq \delta$ holds for some choice of $\delta > 0$.

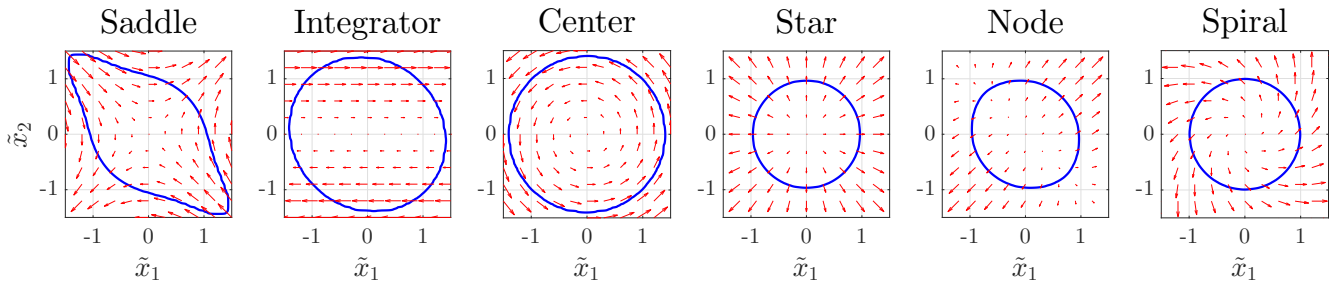

Fig. 6. Optimal sampling policies (blue solid) for six different systems categorized by the phase portrait of their expected trajectories (red arrows). Sampling should be triggered whenever the state leaves the area inside the blue shape.

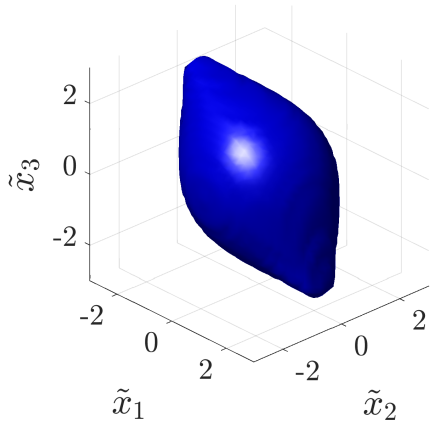

Fig. 7. Example of a non-convex optimal sampling policy in the 3D case, where it is optimal to sample when the error state crosses the surface (blue). Here, $A = [0\,0\,5; 0\,5\,0; 5\,0\,0]$, $R = Q = I$ and $J = 4$. The RBF approximation was generated using $30^3 = 27,000$ collocation points.

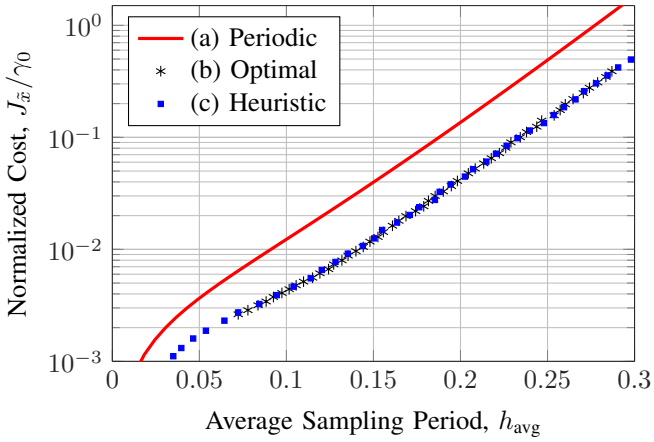

Fig. 8. Trade-off curve between the relative performance degradation $J_{\tilde{x}}/\gamma_0$ and average sampling period h_{avg}. While periodic sampling (red solid) is clearly outperformed by event-based sampling, the difference is negligible between sampling based on the RBF approximation of the optimal policy (black asterisks) and the simple heuristic policy (blue squares).

We consider a 2D unstable system with the following parameters:

$$A = \begin{bmatrix} 0 & 15 \\ 15 & 0 \end{bmatrix}, \quad B_w = C_z^\intercal = \begin{bmatrix} 3.35 & 0 & 0 & 0 \\ -3.27 & 0.72 & 0 & 0 \end{bmatrix},$$

$$B_u = C_y^\intercal = \begin{bmatrix} 28.71 & 0 \\ 28.64 & 2 \end{bmatrix}, \qquad D_{zu} = D_{yw}^\intercal = \begin{bmatrix} 0 \\ I \end{bmatrix},$$

which correspond to $Q = R = I$ and the continuous-time LQG cost $\gamma_0 = 16.41$. Note that for this system the optimal sampling policy will have the same shape as in the saddle point example in Fig. 6.

The trade-off curve for (a) is computed using the expression in [14, Remark 4], while for (b) and (c) we employ Monte-Carlo methods. Specifically, trade-off curves for (b) and (c) are obtained by using different values of J and δ respectively, and simulating a sampled version of the system. The system is sampled with a nominal period $h_{\text{nom}} = 10^{-4}$, and the simulation runs until the standard deviation of the Monte-Carlo estimates of h_{avg} is smaller than 10^{-3}.

The trade-off curves are presented in Fig. 8, where the cost $J_{\tilde{x}}$ has been normalized by γ_0. Note that periodic sampling is clearly outperformed by the event-based sampling policies, where the improvement grows with h_{avg} (note the logarithmic scale in Fig 8). For example, at $h_{\text{avg}} = 0.28$ the closed-loop LQG-cost is increased by roughly 100% for periodic sampling compared to the continuous-time LQG cost, while the corresponding increase is only about 30% for the event-based policies. Secondly, we note that the performance of the simple heuristic policy (c) and the approximation of the optimal policy (b) are practically identical. This further supports the hypothesis that simple, elliptic, policies can achieve near-optimal performance, even for systems with a non-convex optimal policy.

VI. CONCLUSIONS

In this paper we have derived a numerical method using radial basis functions to approximate the optimal event-based sampling policy in the LQG setting. The optimal sampling problem is equivalent to solving a stationary free boundary PDE, and the proposed method is able to solve this PDE efficiently for systems up to order 3. Higher order systems will require further research into efficient distributions of collocation points in order to keep the problem size manageable.

Guarantees for existence and uniqueness of the optimal RBF approximation have been given, and we have shown that it is straightforwardly obtained by solving a QP. The method has been validated numerically, and is shown to converge to the true solution with increasing number of basis functions.

Using the RBF approximation method, we have characterized different types of possible optimal policies in the 2D case. Since most of the policies are convex and almost elliptic, this suggests that simpler, elliptic policies could be designed to achieve near-optimal performance. This was further supported by a numerical example, where the performance of a simple elliptic policy was practically identical to the performance of the optimal one. Future work will be focused on design rules for such sub-optimal policies, which

can be benchmarked against the corresponding optimal ones obtained from the proposed RBF method.

ACKNOWLEDGMENT

The author would like to thank Anton Cervin and Bo Bernhardsson for valuable comments and feedback.

APPENDIX

Proof of Theorem 2

We prove Theorem 2 by showing that $M = -\Lambda\Phi^{-1}$ is a P-matrix (a matrix with positive principal minors) for some finite choice of $c > 0$, since this implies that (24) will have a guaranteed unique solution [20]. To this end we define:

Definition 1: A matrix $B \in \mathbb{R}^{N\times N}$ is *row diagonally dominant* (RDD) if,

$$|B_{i,i}| > \sum_{j\neq i} |B_{i,j}|, \quad \forall i = 1, ..., N,$$

where $B_{i,j}$ denotes the ith row and jth column of B.

We also introduce the following lemma:

Lemma 1 ([23], Prop. 4.6): $-\Lambda\Phi^{-1}$ is a P-matrix if $-\Lambda$ and Φ are RDD matrices with positive diagonal entries.

For proof of this lemma we refer to [23].

First consider Φ, which clearly has positive diagonal entries since $\Phi_{i,i} = \phi_i(\tilde{x}_i) = 1 > 0$, $\forall i$. Thus, for Φ to be RDD, we require that

$$1 > \sum_{j\neq i} |\Phi_{i,j}| = \sum_{j\neq i} \exp(-c\|\tilde{x}_i - \tilde{x}_j\|_2^2), \forall i = 1, ..., N. \tag{27}$$

Note that for $c > 0$, the sum in (27) is a continuous, strictly decreasing function in c. It has the upper and lower limits $N - 1$ and 0 as $c \downarrow 0$ and $c \to \infty$ respectively. Therefore, the inequality (27) is trivially satisfied for $N \leq 2$, while for $N > 2$ there exists a $\underline{c}_{\Phi,i} > 0$ for each row i such that

$$1 = \sum_{j\neq i} \exp(-\underline{c}_{\Phi,i}\|\tilde{x}_i - \tilde{x}_j\|_2^2), \quad \forall i = 1, ..., N.$$

Thus Φ is RDD for $c > \underline{c}_\Phi \triangleq \max_i(\underline{c}_{\Phi,i})$.

Now consider $-\Lambda$, which also has positive diagonal entries since $-\Lambda_{i,i} = nc > 0$, $\forall i$. To be RDD $-\Lambda$ must satisfy

$$\begin{aligned} n &> \frac{1}{c}\sum_{j\neq i} |-\Lambda_{i,j}| \\ &= \sum_{j\neq i} |\tilde{x}_i^\intercal A^\intercal(\tilde{x}_i - \tilde{x}_j) - 2c\|\tilde{x}_i - \tilde{x}_j\|_2^2 + n|\phi_j(\tilde{x}_i), \end{aligned} \tag{28}$$

$\forall i = 1, ..., N.$

Using the triangle inequality we note that the sum in (28) is less than or equal to

$$\sum_{j\neq i} (|\tilde{x}_i^\intercal A^\intercal(\tilde{x}_i - \tilde{x}_j)| + 2c\|\tilde{x}_i - \tilde{x}_j\|_2^2 + n)\phi_j(\tilde{x}_i). \tag{29}$$

Showing that (29) is strictly smaller than n for all rows i is thus sufficient to ensure that (28) is satisfied. The sum (29) is a continuous function in c with the limit 0 as $c \to \infty$, and is guaranteed to be strictly decreasing in c for

$$c > \underline{c}^*_{-\Lambda,i} \triangleq \max_j \left(\frac{2 - |\tilde{x}_i^\intercal A^\intercal(\tilde{x}_i - \tilde{x}_j)| - n}{2\|\tilde{x}_i - \tilde{x}_j\|_2^2} \right).$$

Thus, for every row i we can pick a $\underline{c}_{-\Lambda,i} \geq \underline{c}^*_{-\Lambda,i}$ such that

$$n \geq \sum_{j\neq i} (|\tilde{x}_i^\intercal A^\intercal(\tilde{x}_i - \tilde{x}_j)| + 2\underline{c}_{-\Lambda,i}\|\tilde{x}_i - \tilde{x}_j\|_2^2 + n)\phi_j(\tilde{x}_i).$$

This means that (28) is satisfied for $c > \underline{c}_{-\Lambda} \triangleq \max_i(\underline{c}_{-\Lambda,i})$, implying that $-\Lambda$ is RDD. Finally, if we pick $c > \underline{c} \triangleq \max\{\underline{c}_\Phi, \underline{c}_{-\Lambda}\}$, then both Φ and $-\Lambda$ are RDD, and $-\Lambda\Phi^{-1}$ is a P-matrix according to Lemma 1.

REFERENCES

[1] K.-E. Årzén, "A simple event-based PID controller," *IFAC Proceedings Volumes*, vol. 32, no. 2, pp. 8687–8692, 1999.

[2] K. J. Åström and B. Bernhardsson, "Comparison of periodic and event based sampling for first-order stochastic systems," *IFAC Proceedings Volumes*, vol. 32, no. 2, pp. 5006–5011, 1999.

[3] K. J. Åström, "Event based control," in *Analysis and Design of Nonlinear Control Systems*, A. Astolfi and L. Marconi, Eds. Berlin, Heidelberg: Springer, 2008, pp. 127–147.

[4] W. P. M. H. Heemels, K. H. Johansson, and P. Tabuada, "An introduction to event-triggered and self-triggered control," in *Proc. 51st IEEE Conf. on Decision and Control*, Maui, HI, USA, 2012, pp. 3270–3285.

[5] Q. Liu, Z. Wang, X. He, and D. Zhou, "A survey of event-based strategies on control and estimation," *Systems Science & Control Engineering*, vol. 2, no. 1, pp. 90–97, 2014.

[6] M. Miskowicz, *Event-Based Control and Signal Processing*, ser. Embedded Systems. Boca Raton, FL: CRC Press, 2015.

[7] D. Antunes, W. P. M. H. Heemels, and P. Tabuada, "Dynamic programming formulation of periodic event-triggered control: Performance guarantees and co-design," in *Proc. 51st IEEE Conf. on Decision and Control*, Maui, HI, USA, 2012, pp. 7212–7217.

[8] J. Araújo, A. Teixeira, E. Henriksson, and K. H. Johansson, "A down-sampled controller to reduce network usage with guaranteed closed-loop performance," in *Proc. 53rd IEEE Conf. on Decision and Control*, Los Angeles, CA, USA, 2014, pp. 6849–6856.

[9] D. J. Antunes and B. A. Khashooei, "Consistent event-triggered methods for linear quadratic control," in *Proc. 55th IEEE Conf. on Decision and Control*, Las Vegas, NV, USA, 2016, pp. 1358–1363.

[10] K. Gatsis, A. Ribeiro, and G. J. Pappas, "State-based communication design for wireless control systems," in *Proc. 55th IEEE Conf. on Decision and Control*, Las Vegas, NV, USA, 2016, pp. 129–134.

[11] D. Baumann, J. Zhu, G. Martius, and S. Trimpe, "Deep reinforcement learning for event-triggered control," in *Proc. 57th IEEE Conf. on Decision and Control*, Miami, FL, USA, 2018, pp. 943–950.

[12] F. Solowjow, D. Baumann, J. Garcke, and S. Trimpe, "Event-triggered learning for resource-efficient networked control," in *Proc. of the American Control Conf.*, Milwuakee, WI, USA, 2018, pp. 6506–6512.

[13] L. Mirkin, "Intermittent redesign of analog controllers via the Youla parameter," *IEEE Trans. Autom. Control*, vol. 62, no. 4, pp. 1838–1851, 2017.

[14] M. Braksmayer and L. Mirkin, "H_2 optimization under intermittent sampling and its application to event-triggered control," in *Proc. 20th IFAC World Congress*, Toulouse, France, 2017.

[15] A. Goldenshluger and L. Mirkin, "On minimum-variance event-triggered control," *IEEE Control Syst. Lett.*, vol. 1, no. 1, pp. 32–37, 2017.

[16] M. T. Andrén, B. Bernhardsson, A. Cervin, and K. Soltesz, "On event-based sampling for LQG-optimal control," in *Proc. 56th IEEE Conf. on Decision and Control*, Melbourne, Australia, 2017, pp. 5438–5444.

[17] T. Henningsson, "Stochastic event-based control and estimation," Ph.D. dissertation, Dept. of Automatic Control, Lund University, Sweden, 2012.

[18] B. Fornberg and N. Flyer, "Solving PDEs with radial basis functions," *Acta Numerica*, vol. 24, pp. 215–258, 2015.

[19] M. D. Marcozzi, S. Choi, and C. S. Chen, "On the use of boundary conditions for variational formulations arising in financial mathematics," *Applied Mathematics and Computation*, vol. 124, no. 2, pp. 197–214, 2001.

[20] K. G. Murty, "On the number of solutions to the complementarity problem and spanning properties of complementary cones," *Linear Algebra and its Applications*, vol. 5, no. 1, pp. 65–108, 1972.

[21] K. Zhou, J. C. Doyle, and K. Glover, *Robust and Optimal Control*. Englewood Cliffs, NJ: Prentice-Hall, 1996.

[22] S. M. Rump, "On P-matrices," *Linear Algebra and its Applications*, vol. 363, pp. 237–250, 2003.

[23] M. J. Tsatsomeros, "Generating and detecting matrices with positive principal minors," *Asian Information-Science-Life*, vol. 1, no. 2, pp. 115–132, 2002.