IEEE Control Systems Letters paper presented at
2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024. MiCo, Milan, Italy

Rate-Distortion Achievability via Event Threshold Quantizers for Planar Wiener Processes

Ronald Ogden Takashi Tanaka

***Abstract*— We analyze the rate-distortion performance of two quantized event-based encoding paradigms estimating a two-dimensional Wiener process. Each encoder remains silent until the estimation error reaches a threshold and then transmits a packet from a finite codebook over a noiseless, zero-delay channel to a decoder that updates the estimate. Both encoding methods are parameterized by the radius of the event threshold and the size of the codebook. The first encoding scheme simply quantizes the error of the source process on the event threshold. The second scheme employs a dithered quantizer, which simplifies the derivation of an analytical rate-distortion upper bound. Each method is simulated in discrete time to inform the choice of the bitrate-optimal codebook size and event radius for a given distortion constraint. The rate-distortion performance of these encoding schemes are compared to a known lower bound and a conservative upper bound that we derive herein.**

I. Introduction

Currently, periodic sampling paradigms dominate the fields of sensing, estimation and control theory. However, event-based sampling schemes have shown the potential to reduce the sampling rate required to achieve a fixed estimation quality [1], among other benefits. Unlike periodic methodologies which generate samples at regular time intervals, event-based schemes produce samples as needed according to some application-specific rule. For example, unlike frame-based cameras which encode information for each pixel at each time step, event cameras encode information from pixels asynchronously, each pixel firing only when it detects a significant change in intensity. Event-based sampling methodologies like this enable information to be encoded in both the samples and their timing which can reduce the data rate required to stabilize a system [2]. In practice, this reduction allows sensors, estimators and encoding policies that utilize event-based architectures to achieve a desired performance objective while reducing the consumption of physical resources such as power, computational capacity and bandwidth.

In this work, we explore the performance of event-based sampling when estimating a process over a channel with communication constraints. In such a scenario, one often desires to minimize the communication resources (data rate) required to achieve a given estimation quality (distortion) or vice versa. Demonstrating optimality of an event-based sampler observing a general process has proven to be a challenging problem, but some success has been found for particular processes. For example, optimal sampling of continuous-time 1-dimensional stochastic processes has been studied extensively. Guo and Kostina considered Markov processes in [3], Rabi et al. considered Ornstein-Uhlenbeck processes in [4], Sun et al. considered a Wiener process in [5], and Ornee and Sun considered a Ornstein-Uhlenbeck processes in [6], the latter two considering channel delay. All find that the resulting optimal sampling policy is a threshold-based policy similar to that studied by Åström and Bernhardsson in [1]. The optimality of threshold-based sampling policies for the estimation of discrete time processes has also been extensively studied, as in [7] and [8]. Building on the results of Nayyar et al. in [9], which derive a jointly optimal communication schedule and estimation strategy for remote estimation of a discrete time process, Nar and Başar demonstrate in [10] that for an n-dimensional Wiener process, the real-valued sampling rule that minimizes the time-averaged mean square estimation error given a constraint on average sampling frequency is an analogous event-triggered threshold policy.

In practice, one is often limited to the use of a finite codebook to encode the estimation error of a process, precluding the use of real-valued sampling. Despite this, the problem of finding an optimal sampling policy for the system described in [10] under a finite codebook constraint remains unexplored. For such a problem, it is natural to inquire whether a quantized threshold sampling policy analogous to the optimal policy presented in [10] remains optimal. In this work, we explore the performance of the following analogous sampling policy for a multidimensional Wiener process: N points are distributed around a sphere of radius r and whenever the magnitude of the estimation error reaches the threshold, r, a quantized sample of the state is taken according to which of the N points the estimation error is closest to. Unlike real-valued sampling, a quantized sampling policy like this will result in a non-zero posterior estimation error after each sample almost surely.

Deriving the analytical rate-distortion performance for this sampling scheme poses the challenge that it is not generally possible to symmetrically distribute N points on a sphere in n-dimensional space. However, for a sphere in 2-dimensional space, this is possible for all N. Furthermore, the resulting threshold is a 1-dimensional manifold, allowing an easy application of the dithering scheme utilized in [11], [12], and [13], which simplifies rate-distortion analysis. We restrict this work to a 2-dimensional Wiener process to leverage these advantages.

R. Ogden and T. Tanaka are with the Department of Aerospace Engineering and Engineering Mechanics at the University of Texas at Austin, Austin, TX 78712 USA ronnieogden95@utexas.edu, ttanaka@utexas.edu. This work is supported by NSF Award 1944318.

Introducing dither enables us to derive an expression for the average sampling rate of this scheme, which we use to develop an upper bound on its rate-distortion performance. In contrast, Cuvelier et al. derive a rate-distortion lower bound in [14] for event-based encoding of Gauss-Markov processes, a very general class of processes compared to the 2-dimensional Brownian motion considered herein. Since Brownian motion is a Gauss-Markov process, the performance of the optimal sampling scheme for a 2-dimensional Wiener process with a finite codebook constraint must lie in between the lower bound from [14] and the upper bound we introduce. Finally, we use simulation to determine what N and r yield the optimal rate-distortion performance for the quantized sampling policies we introduce, providing insight for future investigation into the optimal sampling policy for this problem.

II. Problem Formulation

Consider a continuous-time source process $x_t \in \mathbb{R}^2$ that consists of pure 2-dimensional Brownian motion which we seek to estimate with a given distortion over a noiseless, zero-delay communication channel using minimal data rate. An encoder noiselessly observes the source process and can transmit a codeword from a finite codebook over the channel to a decoder as shown in Fig. 1. The decoder then decodes the message sent from the encoder to update its estimate of the source process $\hat{x}_t$. Note that because the channel is noiseless and zero-delay, the encoder knows all of the information that is received by the decoder and thus has access to $\hat{x}_t$ at the same instant the decoder does. The encoder compares the output of the source process x_t to its estimate $\hat{x}_t$ in continuous time, yielding the estimation error $e_t = x_t - \hat{x}_t$ (we assume $e_0 = 0$), and sends a non-empty, variable-length packet from a codebook $\mathcal{C}_N$ to the decoder exactly when the magnitude of the error reaches a threshold, r. Otherwise, the encoder remains silent. In this model, we assume the decoder decodes each codeword individually, precluding the need for an instantaneous code; a nonsingular code is sufficient. The codebook $\mathcal{C}_N$ consists of the N shortest length, non-empty binary strings, that is, the first N elements of the sequence $\{0, 1, 00, 01, 10, 11, 000, ...\}$. Upon each transmission, the decoder updates the estimate $\hat{x}_t$ based on the a priori estimate and the most recent message received. Note that with this formulation, any encoder/decoder pair can be parameterized by the tuple (N, r).

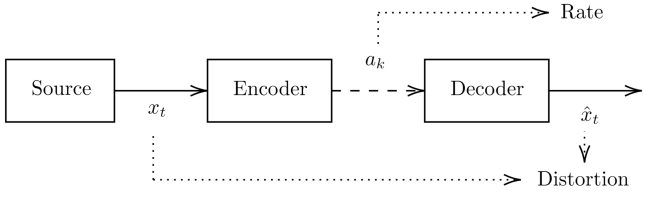

Fig. 1. Block diagram of source process estimation over a channel.

Under these dynamics, one can easily construct an encoder/decoder policy that precludes Zeno behavior; a policy that ensures the posterior error is less than r almost surely is sufficient. In this class of policies, one can represent the transmitted codewords as a sequence, $\{a_k\}$, where $k \in \mathbb{N}$, and one can represent the corresponding sequence of transmission times as $\{t_k\}$. Given an encoder/decoder policy $\pi(N, r)$, we quantify the associated communication cost as the expected bits transmitted per unit of time.

$$R(\pi(N,r)) = \limsup_{T\to\infty} \frac{1}{T} \sum_{k:t_k \le T} \mathbb{E}\left[l(a_k)\right], \tag{1}$$

where $l : \mathcal{C}_N \to \mathbb{N}$ maps a codeword to its length. We quantify the corresponding distortion as the time averaged mean squared error.

$$D(\pi(N,r)) = \limsup_{T\to\infty} \frac{1}{T} \int_0^T \mathbb{E}\left[\|e_t\|_2^2\right] dt. \tag{2}$$

For a fixed policy π, we wish to find the parameterizing tuple (N, r) that minimizes the expected bitrate while achieving a distortion less than D_0. That is, we want to compute

$$\begin{aligned}(N_\pi^*, r_\pi^*) = &\arg \inf_{(N,r)} R(\pi(N,r)) \\ &\text{s.t. } D(\pi(N,r)) \le D_0.\end{aligned} \tag{3}$$

III. Achievability Result

In this section, we describe two policies that fall under the class described in Section II. For this 2-dimensional source process, whenever the magnitude of the estimator error reaches the event threshold of the encoder , the state of the process can be encoded solely with the angular argument $\theta \in \Theta := [0, 2\pi)$ of its polar representation, a fact that is leveraged under both policies described below. Analysis of the dithered quantizer defined in Section III-B shall be used to derive a rate-distortion upper bound in Section III-C.

A. Escape Argument Quantizer

Let Q_Δ be a uniform scalar quantizer defined as follows: $Q_\Delta(\theta) := \Delta\left(\lfloor \theta/\Delta \rfloor + 1/2\right)$, where $\Delta \in \mathbb{R}$ is the quantizer step size. Then the image of the set Θ under the function $Q_{\frac{2\pi}{N}}$, denoted $Q_{\frac{2\pi}{N}}(\Theta)$, has N elements. Thus, if the argument of the source process at an event is θ_{t_k} and $\hat{\theta}_k = Q_{\frac{2\pi}{N}}(\theta_{t_k})$, then $\hat{\theta}_k$ can be losslessly encoded with a codebook of size N. Without loss of generality, we choose any bijection from $Q_{\frac{2\pi}{N}}(\Theta)$ to $\mathcal{C}_N$ to be the encoding policy for an event. The decoder then applies the inverse mapping to the received codeword and updates the estimate in Cartesian coordinates as follows:

$$\hat{x}_{t_k^+} = \hat{x}_{t_k^-} + r \begin{pmatrix} \cos\hat{\theta}_k \\ \sin\hat{\theta}_k \end{pmatrix}. \tag{4}$$

The behavior of this policy is depicted in Fig. 2.

Under this policy, the maximum magnitude of the posterior error is given by $\max_{\theta_{t_k}} \|e_{t_k^+}\| = r\sqrt{2(1-\cos(\pi/N))}$, for given an arc with angle $2\pi/N$, the Euclidean distance between the center of the arc and either edge is this quantity. As mentioned in Section II, we desire this quantity to be less than r almost surely, which holds if we restrict $N \ge 3$. As N approaches infinity, the behavior of this policy approaches

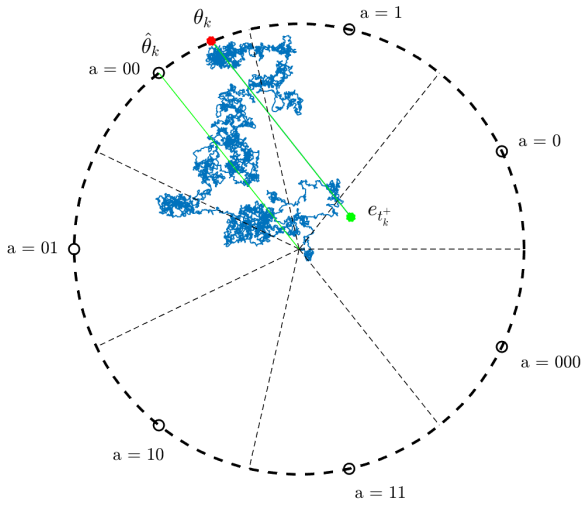

Fig. 2. Illustration of an escape argument quantizer policy. The blue path depicts the evolution of the error e_t between two events with the center of the circle corresponding to an error of 0. When the magnitude of the error reaches the event threshold r, depicted by the black dashed circle, the argument θ_k of the a priori error is quantized to produce $\hat{\theta}_k$. The small black circles on the event threshold represent the possible values of $\hat{\theta}_k$, and in this example, since the escape argument is closest to $a = 00$, this is the codeword that is used to encode the error. The parallel green lines illustrate the direction of the corresponding estimate update.

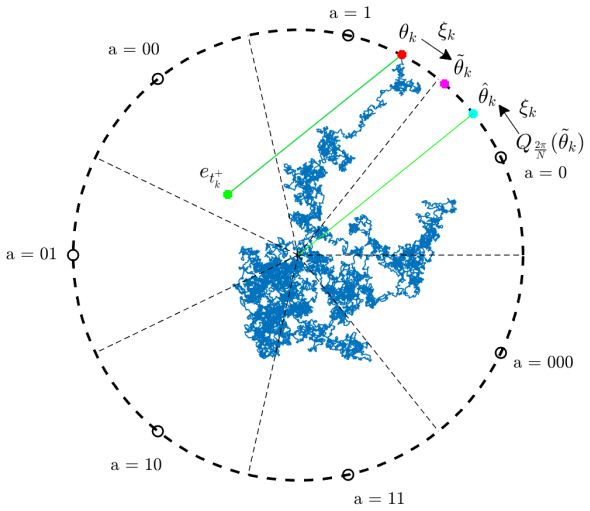

Fig. 3. Illustration of a dithered argument quantizer policy. When the magnitude of the error reaches the event threshold, some dither ξ_k is added to the argument θ_k of the a priori error to produce the dithered escape argument $\tilde{\theta}_k$. This is then quantized to produce $Q_{2\pi/N}(\tilde{\theta}_k)$. The small black circles on the event threshold represent the possible values of $Q_{2\pi/N}(\tilde{\theta}_k)$, and in this example, since the dithered escape argument is closest to $a = 0$, this is the codeword that is used to encode the error. The decoder decodes this transmission and subtracts the same dither ξ_k from the result to produce $\hat{\theta}_k$, which is used in the estimation update. The parallel green lines illustrate the direction of the corresponding estimate update.

that of the real-valued sampling scheme that was shown to be optimal in [10]. The policy presented in this section acts as a quantized analog of the optimal policy proposed in [10], making it a natural policy to examine given a finite codebook constraint. However, the derivation of an analytical rate-distortion performance can be simplified with the introduction of dither.

B. Dithered Argument Quantizer

This policy, in short, is a dithered version of the policy described in Section III-A. At the k^{th} event, the encoder and decoder are given access to the value of the same uniform random variable $\xi_k \sim \mathcal{U}[-\frac{\pi}{N}, \frac{\pi}{N}]$, known as the dither (in practice, this could be achieved by using synchronized pseudorandom number generators at each end of the communication channel). Suppose the argument of the source process at an event is θ_{t_k}. Then let $\tilde{\theta}_k := (\theta_{t_k} + \xi_k) \bmod 2\pi \in \Theta$ be the dithered escape argument. As before, $Q_{\frac{2\pi}{N}}(\tilde{\theta}_k)$ can be losslessly encoded using any bijection from $Q_{\frac{2\pi}{N}}(\Theta)$ to $\mathcal{C}_N$. The decoder then recovers this quantity using the inverse of this mapping and proceeds to compute $\hat{\theta}_k = Q_{\frac{2\pi}{N}}(\tilde{\theta}_k) - \xi_k$ and updates the estimate of the process according to (4). The behavior of this policy is depicted in Fig. 3

Let $\eta_k := Q_{\frac{2\pi}{N}}(\tilde{\theta}_k) - \tilde{\theta}_k$ be the quantization error of the dithered argument. Note that $\eta_k \sim \mathcal{U}[-\frac{\pi}{N}, \frac{\pi}{N}]$ and is independent of θ_{t_k}. This fact has been demonstrated in several previous works including [12], [13] and [15]. Furthermore, $\hat{\theta}_k \equiv (\eta_k + \theta_{t_k}) \bmod 2\pi$, so under this policy, $\hat{\theta}_k$ is equivalent to the escape argument plus uniform noise. This means that at every event, the distribution of the magnitude of the posterior estimation error is known, which we shall leverage to facilitate the analysis of this policy. As with the undithered encoder, we restrict $N \geq 3$ to ensure the posterior error is almost surely within the event threshold.

C. Rate-distortion Upper Bound

We combine a precise computation of the rate achievable by the dithered quantizer with a conservative upper bound on its distortion that applies to all policies described in Section II. The result is the following upper bound on the performance achievable by a dithered quantizer.

Theorem 1: Let $\pi_{\mathcal{D}}$ denote the dithered policy described in Section III-B. Suppose $D(\pi_{\mathcal{D}}(N, r)) = D_0$. Then the following holds:

$$R(\pi_{\mathcal{D}}(N, r)) \leq \frac{\frac{1}{N}\sum_{i=1}^{N} \lfloor \log_2(i+1) \rfloor}{D_0\left(\frac{2N}{\pi}\sin\left(\frac{\pi}{N}\right) - 1\right)}.$$

Proof: Let $\tau_1 := t_1$ and $\tau_k := t_k - t_{k-1}$ for $k > 1$. The communication cost as defined in (1) of this policy is given by

$$\begin{aligned} R(\pi_{\mathcal{D}}(N, r)) &= \limsup_{K\to\infty} \left(\sum_{k=1}^{K} \tau_k\right)^{-1} \sum_{k=1}^{K} \mathbb{E}\left[l(a_k)\right] \\ &= \limsup_{K\to\infty} \left(\frac{\sum_{k=1}^{K} \tau_k}{K}\right)^{-1} \mathbb{E}\left[l(a_k)\right] \\ &= \frac{\mathbb{E}\left[l(a_k)\right]}{\mathbb{E}\left[\tau_k\right]}, \end{aligned} \tag{5}$$

where the first equality comes from the fact that $T - t_K$ is small compared to t_K as $T \to \infty$, the second equality is due to the radial symmetry of both the policy and the dynamics

(no codeword is used preferentially), and the final equality is the result of the law of large numbers.

To compute $\mathbb{E}[l(a_k)]$, we begin by considering the first K codewords sent by the encoder $\{a_k\}_{k=1}^K$. Let $f : \mathcal{C}_N \mapsto \mathcal{C}_N$ be an arbitrary bijection of $\mathcal{C}_N$ to itself. The initial estimation error is 0 and independent of any information, the distribution of x_t has polar symmetry. Therefore any sequence $\{a_k\}_{k=1}^K$ is as likely to occur as $\{f(a_k)\}_{k=1}^K$. The bijection f is arbitrary, so the expected frequency of any codeword is uniform. Thus, the expected length of the codewords used is the average length of the available codewords in $\mathcal{C}_N$. The length of the i^{th} element of the sequence $\{0, 1, 00, 01, 10, 11, 000, ...\}$ is $\lfloor \log_2(i+1) \rfloor$, so we have

$$\mathbb{E}[l(a_k)] = \frac{1}{N} \sum_{i=1}^{N} \lfloor \log_2(i+1) \rfloor . \tag{6}$$

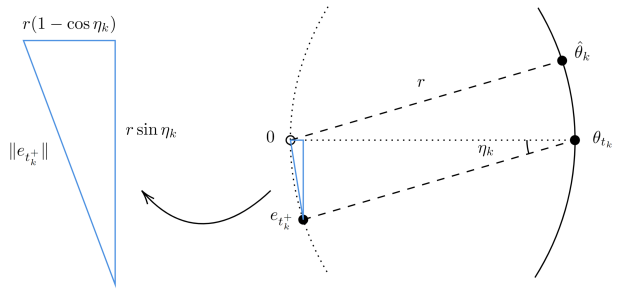

Fig. 4. Geometric illustration of a dithered quantizer estimation update. The solid arc depicts a section of the event threshold, the parallel dashed lines indicate the direction of the estimation update and the dotted arc represents the possible values the posterior error can take given the escape argument is θ_{t_k}. The distribution of the posterior error given the escape argument is uniform along this arc.

To compute $\mathbb{E}[\tau_k]$, we begin by noting that by Theorem 9.1.2 in [16], the expected escape time of a Wiener process from a d-dimensional ball of radius r starting from a point e inside the ball is given by

$$\mathbb{E}[\tau] = \frac{r^2 - \|e\|^2}{d}. \tag{7}$$

Given the quantization error for the k^{th} event, η_k, the magnitude of the posterior estimation error is given by $\|e_{t_k^+}(\eta_k)\| = r\sqrt{2(1-\cos\eta_k)}$. See Fig. 4 for a geometric illustration. Additionally, for any event, $\eta_k \sim \mathcal{U}[-\frac{\pi}{N}, \frac{\pi}{N}]$. Combining these facts with (7) and utilizing the law of iterated expectations, we have

$$\begin{aligned}\mathbb{E}[\tau_k] &= \mathbb{E}[\mathbb{E}[\tau_k|\eta_k]] = \int \mathbb{E}[\tau_k|\eta_k = \eta]\, p(\eta) d\eta \\ &= \int_{-\pi/N}^{\pi/N} \frac{r^2 - 2r^2(1-\cos\eta)}{2} \left(\frac{N}{2\pi}\right) d\eta \\ &= r^2 \left(\frac{N}{\pi} \sin\left(\frac{\pi}{N}\right) - \frac{1}{2}\right). \end{aligned} \tag{8}$$

As expected, as N approaches infinity, (8) shows that the sampling frequency of the dithered quantizer approaches the behavior of the real-valued sampling analog studied in [10]. Substituting (6) and (8) into (5) yields an expression for the bitrate of $\pi_{\mathcal{D}}$.

We provide a conservative upper bound for the distortion of $\pi_{\mathcal{D}}$. Let ρ represent the magnitude of the posterior error. Then as before, we have $\rho = r\sqrt{2(1-\cos\eta)}$. Given that $\eta \sim \mathcal{U}[-\frac{\pi}{N}, \frac{\pi}{N}]$, the probability that the posterior error is less than ρ is the fraction of values of η in the interval $[-\frac{\pi}{N}, \frac{\pi}{N}]$ that correspond to a posterior error of less than ρ. (In Fig. 4, it is the fraction of the dotted arc that is within a distance ρ of the origin). Thus, letting $\rho_{\max} := r\sqrt{2(1-\cos(\pi/N))}$ be the magnitude of the maximum posterior error,

$$\mathbb{P}\left(\|e_{t_k^+}\| \le \rho\right) = \begin{cases} \frac{N}{\pi} \cos^{-1}\left(1 - \frac{1}{2}\left(\frac{\rho}{r}\right)^2\right), & \rho < \rho_{\max}, \\ 1, & \rho \ge \rho_{\max}. \end{cases}$$

Differentiating this expression with respect to ρ and dividing by $2\pi\rho$ yields an expression for the polar coordinate representation of the 2-dimensional PDF of the posterior distribution:

$$p(\rho, \theta) = \begin{cases} \frac{N}{2\pi^2 \rho \sqrt{r^2 - \rho^2/4}}, & \rho < \rho_{\max}, \\ 0, & \rho \ge \rho_{\max}. \end{cases} \tag{9}$$

For a fixed θ, the PDF in (9) for the posterior error distribution strictly decreases as ρ increases over the interval $(0, \rho_{\max})$. The process e_t is composed of Brownian diffusion and a jump process from the event boundary to the posterior distribution in (9). Thus, for a fixed θ, the stationary (i.e. steady state) distribution for the estimation error strictly decreases as ρ increases over the interval $(0, r)$, resulting in a stationary distribution that has a variance smaller than that of a uniform distribution on a disk of radius r. The variance of this uniform distribution is given by $r^2/2$, so we have that $D(\pi_{\mathcal{D}}(N, r)) \le r^2/2$. ■

As shown in [10], the distortion of this policy as $N \to \infty$ is given by $r^2/4$, so this bound is conservative. A precise rate-distortion curve could be derived given an explicit expression for the stationary distribution e_t. As our simulations will show, the optimal dithered and non-dithered rate-distortion curves are very close, so such an expression could also be used to accurately gauge the performance of the optimal escape argument quantizer. For the dithered policy, $\hat{\theta}_k$ is equivalent to θ_{t_k} plus uniform noise, so this policy can perform no better than the undithered policy. Thus, Theorem 1 bounds the performance of the escape argument quantizer as well.

IV. Discrete-time Simulation

We simulated both of the policies described in Section III to determine the optimal tuple (N^*, r^*) for each policy and how the rate-distortion performance of the policies compare to each other, Theorem 1 and a lower bound from [14].

A. Simulation Description

In order to conduct numerical experiments, the continuous time process e_t is simulated with a sufficiently small step size Δt at each time step i using the following dynamics: $e[i+1] = e[i] + w[i]\sqrt{\Delta t} - \Delta\hat{x}[i]$, where each $w[i] \sim \mathcal{N}(0, I)$

is i.i.d. and $\Delta\hat{x}[i]$ is the estimation update which takes the value 0 for all i such that $\|e[i]\| < r$ and is assigned the following otherwise: $\Delta\hat{x}[i] = r(\cos\hat{\theta}[i], \sin\hat{\theta}[i])^\top$.

The value of $\hat{\theta}[i]$ depends on the encoder being used. For the escape argument quantizer described in Section III-A, the simulated estimation update is given by $\hat{\theta}[i] = Q_{\frac{2\pi}{N}}(\arg e[i])$. For the dithered argument quantizer described in Section III-B, the simulated estimation update is given by

$$\hat{\theta}[i] = Q_{\frac{2\pi}{N}}\big((\arg e[i] + \xi[i]) \bmod 2\pi\big) - \xi[i],$$

where each $\xi[i] \sim \mathcal{U}[-\frac{\pi}{N}, \frac{\pi}{N}]$ is i.i.d.

To compute the bitrate for the simulation, an arbitrary map from $Q_{\frac{2\pi}{N}}(\Theta)$ to $\mathcal{C}_N$ was selected as in Fig. 2 so that the k^{th} event would generate a codeword a_k from the input to $Q_{\frac{2\pi}{N}}(\cdot)$. Then, assuming the simulation was run over K events, the bitrate was computed as the ratio of the total bits sent between the encoder and the decoder over the simulation divided by the total simulation time: $R_{\text{sim}} = [\sum_{k=1}^{K} l(a_k)]/(i_K \Delta t)$.

The corresponding distortion was computed as the average estimation error squared, $\|e[i]\|^2$, over all timesteps: $D_{\text{sim}} = (\sum_{i=1}^{i_K} \|e[i]\|^2)/i_K$.

B. Simulation Results

In order to determine N^* and r^* for both of the policies described in this paper, the simulation described in Section IV-A was run for each policy with $e[0] = 0$ over values of $N \in \{3, 4, ..., 8\}$ and $r \in \{2, 2.5, ..., 6\}$. The time step of the simulation Δt was scaled with the square of the threshold radius r. In particular, these simulations used $\Delta t = r^2/200$. Each simulation was run until 10,000 events occurred. The results are shown in Fig. 5 and Fig. 6.

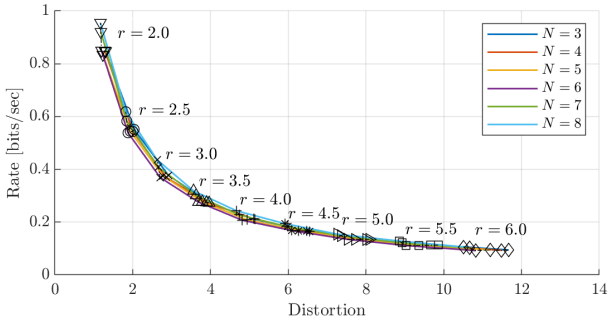

Fig. 5. Simulation results for the escape argument quantizer.

Observe that if one fixes r, the resulting bitrate is not a monotonic function of N. This is because on one hand, the definition of our variable length codebook $\mathcal{C}_N$ makes the expected codeword length a strictly increasing function of N, so increasing N causes the numerator of (5) to increase. However, increasing N reduces the posterior error at each event, resulting in an increase in the expected sampling period, the denominator of (5).

Both Fig. 5 and Fig. 6 show that the rate-distortion curve corresponding to $N = 6$ lies below that of any other codebook size. Therefore, these simulations suggest that for

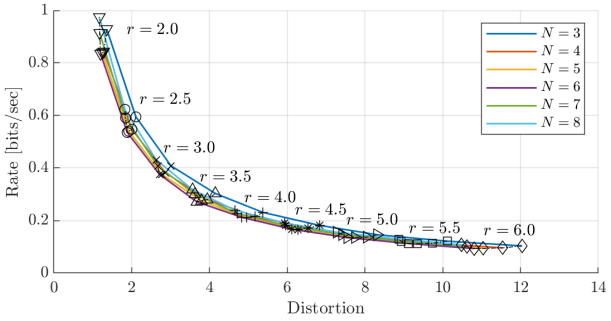

Fig. 6. Simulation results for the dithered argument quantizer.

either of the policies described in this paper, the optimal codebook size for a given distortion constraint is $N^* = 6$. Fixing this parameter, the corresponding optimal threshold r^* is a one-to-one function of the distortion constraint. For example, Fig. 6 shows that for a dithered argument quantizer, given a distortion constraint $D(\pi(N, r)) \leq 6.1$, the optimal codebook size is $N^* = 6$ and the corresponding optimal threshold is approximately $r^* = 4.5$, achieving a bit rate of approximately $R(\pi(N^*, r^*)) = 0.17$ bits/sec.

It is worth noting that N^* depends on the type of codebook used. For example, if we use a fixed-length codebook, simulations show $N^* = 4$, which can be explained in part by the relatively large jump in expected codeword length from 2 for $N \in \{3, 4\}$ to 3 for $N \in \{5, 6, 7, 8\}$ for such a codebook, as shown in Fig. 7.

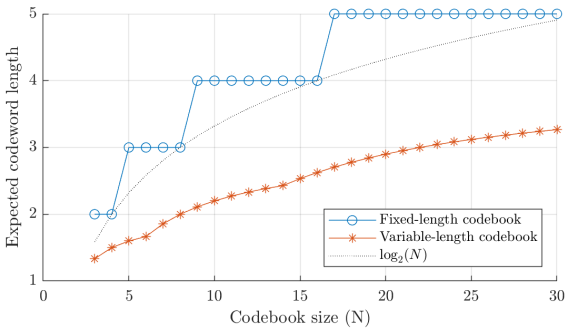

Fig. 7. Comparison of expected codeword length.

While the upper bound derived in Theorem 1 holds for any $N \geq 3$, we seek to compare the optimal version of the encoders presented in this work (i.e. those with $N = 6$) to known lower bounds. The simulation results of both policies for $N = 6$ are shown in Fig. 8 and compared to the upper bound derived in Section III-C and to the continuous time information-distortion (CTID) lower bound given in Theorem 1 of [14]. We state the implication of that result as it pertains to this example for convenience.

Corollary 1: Suppose a sampling policy $\pi(N, r)$ observing a 2-dimensional Weiner process is restricted to observing the process at discrete instances of time: $\{\Delta t, 2\Delta t, 3\Delta t, ...\}$. Suppose additionally that $D(\pi(N, r)) = D_0$. Then

$$R(\pi(N, r)) \geq \frac{1}{\Delta t}\theta^{-1}(\Delta t \mathcal{I}^c(D_0)),$$

where $\theta(x) := x + (1+x)\log_2(1+x) - x\log_2(x)$ and

$$\mathcal{I}^c(D_0) := \inf_{X\geq 0, Y\geq 0} \frac{\mathrm{Tr}(Y)}{2\ln 2}, \quad \text{s.t.} \quad \begin{pmatrix} Y & I \\ I & X \end{pmatrix} \geq 0, \\ \mathrm{Tr}(X) \leq D_0,$$

where I is the 2-dimensional identity matrix.

Note that this lower bound depends on Δt, so the smallest and thus most conservative Δt used across all simulations was selected to plot the lower bound. These results demonstrate that for $N = 6$, the performance of these policies is nearly the same. Although the upper bound from Section III-C appears to be fairly loose, the similarity in performance between both policies suggests that if an analytical expression for distortion of the dithered argument quantizer could be derived, one would have an accurate expression for the rate-distortion curve of that policy and thus an accurate estimate of the rate-distortion performance of the escape argument quantizer.

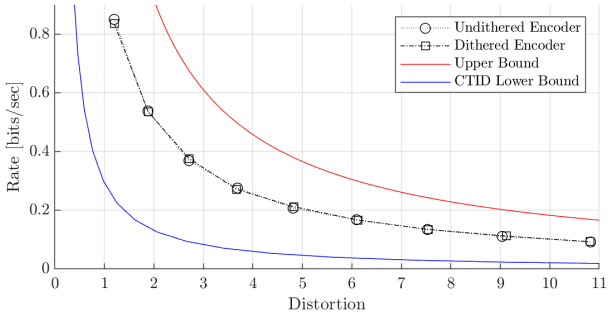

Fig. 8. Rate-distortion comparison of escape argument quantizer and dithered argument quantizer against the upper bound derived in Theorem 1 and the CTID lower bound presented in Corollary 1.

V. Conclusion and Future Work

In this work, we analyzed quantized threshold sampling policies to estimate a Wiener process in two dimensions and found an upper bound on their performance. While this work is limited to an achievability result for a Wiener process in two dimensions, the analysis herein serves as a bridge to analyzing similar sampling policies for richer source processes in higher dimensions.

This work was restricted to two dimensions in part due to the loss of symmetry for N points uniformly distributed around a sphere in higher dimensions. However, in n-dimensional space for $n \geq 2$ there always exists a uniform, symmetric arrangement of N points about a sphere for $N \in \{n+1, 2n, 2^n\}$, as there exist regular polytopes with those vertex counts in those dimensions [17]. We can utilize the symmetry of these special cases to extend the quantized sampling policy presented in this paper to monitor higher dimensional Brownian motion while still maintaining analytical tractability. Furthermore, simulating the performance of a quantized threshold sampling policy in higher dimensions can allow us to compare the performance of this policy to existing rate-distortion lower bounds like those presented in [14]. Spherical codes are well-studied, and databases of optimized spherical codes are readily available, such as [18] which can facilitate the development of simulations in higher dimensions.

Finally, we wish to explore the optimal quantization of a sampling policy with a fixed event threshold when a linear drift term is introduced in the source process. Such an investigation would provide insight into the design of sampling schemes to optimize rate-distortion performance for more complex processes such as the event camera discussed in the introduction.

References

[1] K. J. Astrom and B. M. Bernhardsson, "Comparison of riemann and lebesgue sampling for first order stochastic systems," in *Proceedings of the 41st IEEE Conference on Decision and Control, 2002.*, vol. 2, pp. 2011–2016, IEEE, 2002.

[2] M. J. Khojasteh, P. Tallapragada, J. Cortés, and M. Franceschetti, "The value of timing information in event-triggered control," *IEEE Transactions on Automatic Control*, vol. 65, no. 3, pp. 925–940, 2019.

[3] N. Guo and V. Kostina, "Optimal causal rate-constrained sampling for a class of continuous Markov processes," *IEEE Trans. on Info. Theory*, vol. 67, no. 12, pp. 7876–7890, 2021.

[4] M. Rabi, G. V. Moustakides, and J. S. Baras, "Adaptive sampling for linear state estimation," *SIAM journal on control and optimization*, vol. 50, no. 2, pp. 672–702, 2012.

[5] Y. Sun, Y. Polyanskiy, and E. Uysal-Biyikoglu, "Remote estimation of the wiener process over a channel with random delay," in *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 321–325, IEEE, 2017.

[6] T. Z. Ornee and Y. Sun, "Sampling for remote estimation through queues: Age of information and beyond," in *2019 International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)*, pp. 1–8, IEEE, 2019.

[7] G. M. Lipsa and N. C. Martins, "Remote state estimation with communication costs for first-order lti systems," *IEEE Transactions on Automatic Control*, vol. 56, no. 9, pp. 2013–2025, 2011.

[8] A. S. Leong, S. Dey, and D. E. Quevedo, "Sensor scheduling in variance based event triggered estimation with packet drops," *IEEE Transactions on Automatic Control*, vol. 62, no. 4, pp. 1880–1895, 2016.

[9] A. Nayyar, T. Başar, D. Teneketzis, and V. V. Veeravalli, "Optimal strategies for communication and remote estimation with an energy harvesting sensor," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2246–2260, 2013.

[10] K. Nar and T. Başar, "Sampling multidimensional wiener processes," in *53rd IEEE Conference on Decision and Control*, pp. 3426–3431, IEEE, 2014.

[11] E. I. Silva, M. S. Derpich, and J. Ostergaard, "A framework for control system design subject to average data-rate constraints," *IEEE Transactions on Automatic Control*, vol. 56, no. 8, pp. 1886–1899, 2010.

[12] R. Zamir and M. Feder, "On universal quantization by randomized uniform/lattice quantizers," *IEEE Transactions on Information Theory*, vol. 38, no. 2, pp. 428–436, 1992.

[13] T. Tanaka, K. H. Johansson, T. Oechtering, H. Sandberg, and M. Skoglund, "Rate of prefix-free codes in lqg control systems," in *2016 IEEE International Symposium on Information Theory (ISIT)*, pp. 2399–2403, IEEE, 2016.

[14] T. Cuvelier, R. Ogden, and T. Tanaka, "Minimum bitrate neuromorphic encoding for continuous-time gauss-markov processes," *arXiv preprint arXiv:2309.06504 (Accepted to IEEE Transactions on Automatic Control)*, 2023.

[15] R. Zamir, "How to generate a simple dither," in *2010 IEEE 26-th Convention of Electrical and Electronics Engineers in Israel*, pp. 000844–000848, IEEE, 2010.

[16] R. Durrett, *Probability: theory and examples*, vol. 49. Cambridge university press, 2019.

[17] H. S. M. Coxeter, *Regular polytopes*. Courier Corporation, 1973.

[18] H. Cohn, "Table of spherical codes," 2024.