Near-Coherent QPSK Performance With Coarse Phase Quantization: A Feedback-Based Architecture for Joint Phase/Frequency Synchronization and Demodulation

Aseem Wadhwa and Upamanyu Madhow

***Abstract*—As communication systems scale up in bandwidth, the limited resolution in high-speed analog-to-digital converters (ADCs) is a key challenge in realizing low-cost "mostly digital" transceiver architectures. This motivates a systematic effort to understand the limits of such architectures under the severe quantization constraints imposed by the use of low-precision ADCs. In particular, we investigate a canonical problem of blind carrier phase and frequency synchronization with coarse phase quantization in this paper. We develop a Bayesian approach to blind phase estimation, jointly modeling the unknown data, unknown phase and the quantization nonlinearity. We highlight the crucial role of dither, implemented via a mixed signal architecture with a digitally controlled phase shift prior to the ADC. We show the efficacy of random dither, and then improve upon its performance with a simple feedback control policy that is close to optimal in terms of rapidly reducing the mean squared error of phase estimation. This initial blind phase acquisition stage is followed by feedback-based phase/frequency tracking using an Extended Kalman Filter. Performance evaluations for a QPSK system show that excellent bit error rate (BER) performance, close to that of an *unquantized* system, is achieved by the use of 8 phase bins (implementable using 4 one-bit ADCs operating on linear combinations of in-phase and quadrature components).**

***Index Terms*—Low precision ADC, synchronization, Bayesian estimation, mixed signal architecture, adaptive control, frequency tracking.**

I. Introduction

MODERN communication transceivers (e.g., for WiFi and cellular systems today) are based on a "mostly digital" architecture, using digital signal processing (DSP) to implement sophisticated functionalities such as synchronization, equalization, demodulation and decoding, thus leveraging the economies of scale resulting from Moore's law. The central assumption in such designs is that analog signals can be faithfully represented in the digital domain, typically using high-precision (e.g., 8–12 bits) ADCs. However, the cost and power consumption of high-precision ADCs become prohibitive at multi-GHz sampling rates [1], which raises the question of whether DSP-centric architectures scale as communication bandwidths increases, such as for emerging millimeter wave wireless networks (e.g., using the 7 GHz of unlicensed spectrum in the 60 GHz band), as well as for optical and backplane communication. In particular, it is of fundamental interest to understand the limits of such architectures, and to devise algorithms for attaining them, when ADC precision is severely reduced (e.g., to 1–4 bits).

Shannon-theoretic analysis for idealized channel models has shown that the loss in channel capacity due to limited ADC precision is relatively small even at moderately high signal-to-noise ratios (SNRs) [2]. This motivates a systematic investigation of DSP algorithms for estimating and compensating for channel non-idealities (e.g., asynchronism, dispersion) using severely quantized inputs. The present paper takes a step in this direction by considering a canonical problem of blind carrier phase/frequency synchronization. Our goal is to obtain fundamental insight into the implications of coarse quantization, rather than to provide a complete link design. We therefore do not model timing asynchronism or channel dispersion, and study the simplest setting of coherent reception of QPSK over an AWGN channel. We consider *phase-only quantization*, which suffices for hard decisions with PSK constellations, and has the advantage of not requiring automatic gain control (AGC), since it can be implemented by passing linear combinations of the in-phase (I) and quadrature (Q) components through one-bit ADCs (quantization into $2n$ phase bins requires n such linear combinations). We develop and evaluate the performance of a Bayesian approach based on joint modeling of the unknown data, frequency and phase, and the known quantization nonlinearity, using a mixed signal architecture in which digitally controlled phase shifts are applied to the samples prior to phase quantization.

A. Receiver Architecture

In the model depicted in Fig. 1, the *analog preprocessing front-end* performs downconversion, ideal symbol rate sampling, and applies a digitally controlled *derotation phase* on the

Manuscript received October 13, 2015; revised March 24, 2016; accepted April 16, 2016. Date of publication May 12, 2016; date of current version July 21, 2016. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Joseph Cavallaro. This work was supported in part by the Institute for Collaborative Biotechnologies under Grant W911NF-09-0001 from the U.S. Army Research Office, and in part by the Systems on Nanoscale Information fabriCs (SONIC), one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP), a Semiconductor Research Corporation program sponsored by MARCO and DARPA. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

The authors are with the Department of Electrical and Computer Engineering, University of California Santa Barbara, Santa Barbara, CA 93106 USA (e-mail: aseem@ece.ucsb.edu; madhow@ece.ucsb.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSP.2016.2568169

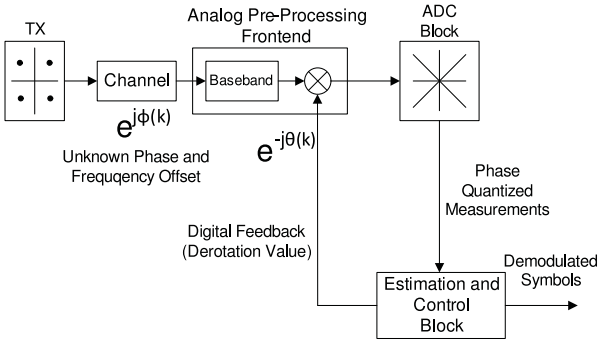

Fig. 1. Receiver Architecture.

complex-valued symbol rate samples before passing it through the *ADC block*. The quantized phase observations are processed in DSP by the *estimation and control block*: this runs algorithms for nonlinear phase and frequency estimation, computes feedback for the analog preprocessor (to aid in estimation and demodulation), and outputs demodulated symbols. The design of this estimation and control block is the subject of this paper.

The frequency offset between transmitter and receiver is typically much smaller than the symbol rate, allowing us to accurately approximate the phase as a constant over a few symbol periods. We can therefore divide the synchronization problem into two stages: (1) rapid blind *acquisition* of initial phase, (2) continuous phase/frequency *tracking* while performing data demodulation. In the tracking stage, the derotation phase is simply an estimate of the (negative of the) overall phase offset. In the acquisition stage, it is not clear *a priori* how to choose the derotation phase. For unquantized (or finely quantized) samples, we could simply set it to zero. However, as we shall see, an appropriate choice of the derotation phase, which serves as a controllable and variable *dither*, is a crucial tool for estimation with severely quantized observations, especially at high SNR. Thus, a significant portion of this paper is dedicated to investigation of dithering strategies for Bayesian phase estimation, including open loop pseudorandom dither as well as feedback control.

B. Contributions

Our contributions are summarized as follows:

1) For the acquisition stage, we develop a Bayesian algorithm for blind phase estimation with coarse phase quantization, and highlight the need for dither. After showing that random open-loop dither works well, we investigate the problem of optimal dither, which falls in the general category of *control for sequential estimation*, finding an exact solution to which is known to be computationally intractable. While several asymptotically optimal policies have been proposed in the literature, these need not be optimal for the small number of measurements of interest to us. We propose a greedy strategy that chooses the feedback to minimize the uncertainty (Shannon entropy) in the posterior distribution of the phase, prove that it converges to an asymptotically optimal policy, while showing via numerical results that it is close to "genie-optimal" for a small number of samples.
2) For the tracking/demodulation stage, we use a two-tier algorithm: decision-directed phase estimation over blocks, ignoring frequency offsets, and an extended Kalman filter (EKF) for long-term frequency/phase tracking. The feedback to the analog preprocessor now aims to compensate for the phase offset, in order to optimize the performance of coherent demodulation with differential decoding. We provide numerical results demonstrating the efficacy of our approach for both steps, and show that the bit error rate with 8–12 phase bins (implementable using linear I/Q processing and 4–6 one bit ADCs) is close to that of a coherent system, and is significantly better than that of standard differential demodulation (which does not require phase/frequency tracking) with unquantized observations.

C. Related Work

A phase-quantized carrier-asynchronous system model similar to ours was studied in [3], but it employs block noncoherent demodulation, approximating the phase as constant over a block of symbols. This approach incurs a loss of about 2 dB with respect to unquantized block noncoherent demodulation, unlike our approach of explicit phase/frequency estimation and compensation, which attains performance almost identical to an unquantized coherent system. A receiver architecture similar to ours (mixed signal analog front-end and low-power ADC with feedback from a DSP block) was implemented for a Gigabit/s 60 GHz system in [4], including blocks for both carrier synchronization and equalization. While the emphasis in [4] was on establishing the feasibility of integrated circuit implementation rather than algorithm design and performance evaluation as in this paper, it makes a compelling case for architectures such as those in Fig. 1 for low-power mixed signal designs at high data rates. Some of the other related work on estimation using low-precision samples includes frequency estimation [5], amplitude estimation for PAM (pulse amplitude modulation) signaling [6], channel estimation [7], equalization [8] and multivariate parameter estimation from dithered quantized data [9]. We postpone discussion of related literature in control for estimation to Section IV, where we describe our greedy feedback control policy and place it in the context of past research.

A preliminary version of this work was presented in a conference paper [10], in which we proposed the information-theoretic greedy control strategy and evaluated its performance via numerical simulations. The present paper goes well beyond [10] in terms of theoretical analysis, as well as more comprehensive performance evaluation for both acquisition and tracking.

The rest of the paper is organized as follows. In Section II, we describe the complex baseband system model. In Sections III and IV, we discuss the procedure of rapid acquisition of an initial estimate of the phase and the control policy for setting the

feedback. In Section V, we present the phase/frequency tracking algorithm and discuss the concluding remarks in Section VI.

II. System Model

We now specify a mathematical model for the receiver architecture depicted in Fig. 1. The analog preprocessor applies a phase derotation of $e^{-j\theta_k}$ for the kth sample. In order to simplify digital control of the derotation, we restrict the allowable *derotation values* θ to a finite set of values. In our simulations, we restrict it to the integer multiples of $\pi/180$ (or 1°). After derotation, the sample is quantized using n 1-bit ADCs into one of $M = 2n$ phase bins: $\left[(m-1)\frac{2\pi}{M}, m\frac{2\pi}{M}\right)$ for $m = 1, ..., M$. In this paper, we consider $M = 8$ and $M = 12$ (Fig. 3(a) and (c)). As mentioned earlier, such phase quantization can be easily implemented by taking n linear combinations of I and Q samples followed by 1-bit ADCs. For example, $M = 8$ bins can be obtained by 1-bit quantization of $I, Q, I+Q$ and $I-Q$. We always include boundaries coinciding with the I and Q axes, since these are the maximum likelihood decision boundaries for coherent QPSK demodulation.

Denoting the phase-quantized observation corresponding to the kth symbol by z_k, we therefore have the following complex baseband measurement model:

$$z_k = Q_M\left(\arg\left(\left(b_k e^{j\phi_k} + w_k\right)e^{-j\theta_k}\right)\right) \in \{1, 2, ..., M\} \quad (1)$$

where b_k are the transmitted QPSK symbols and w_k is complex white Gaussian noise. b_k are uniformly drawn from the set $\left\{e^{j\pi/4}, e^{j3\pi/4}, e^{j5\pi/4}, e^{j7\pi/4}\right\}$. Without loss of generality, we assume magnitude of b_k to be 1 so that $\mathrm{Re}(w_k) = \mathrm{Im}(w_k) \sim \mathcal{N}\left(0, \sigma^2\right)$ where SNR per bit $= \frac{E_b}{N_0} = \frac{1}{2\sigma^2}$.

The time varying phase offset, ϕ_k, depends on the initial offset, ϕ_0 (at time 0) and the frequency offset, Δf.

$$\phi_k = \phi_0 + k\eta T_s; \quad \eta = 2\pi\Delta f \quad (2)$$

T_s is the symbol time period. ηT_s represents the rate of change of phase in radians per symbol. The carrier frequency offset Δf is typically of the order of 10–100 ppm (parts per million) of the carrier frequency, whereas symbol rates are of the order of 1–10% of the carrier frequency, hence ηT_s is small. For example, consider the following typical values: $f_c = 60$ GHz, bandwidth of 6 GHz, i.e., $T_s = \left(6 \times 10^9\right)^{-1}$ secs, an offset $\Delta f = 50$ ppm $\cdot f_c$, which leads to $\eta T_s = 2\pi\Delta f T_s = \pi \cdot 10^{-3}$ radians, or a linear phase change of 0.18° per symbol. Thus, the phase offset is well approximated as constant over a few tens of symbols. This allows us to break the problem into a rapid phase acquisition stage assuming zero frequency offset (Sections III and IV), followed by decoding and tracking initialized with the phase estimate of the first stage (Section V). We assume that the latter recovers the phase modulo $\pi/2$, hence we employ coherent demodulation followed by differential decoding across consecutive symbols. This incurs at most a factor of two degradation in symbol error rate with respect to coherent demodulation with per-symbol absolute decoding (a negligible degradation in dB at even moderate SNRs). Thus, our explicit estimation and compensation strategy (with severe quantization) performs significantly better than two-symbol differential demodulation even with unquantized observations. Block noncoherent demodulation with unquantized observations is known to approach coherent performance as block size increases [11], but as noted earlier, block noncoherent demodulation with severe phase quantization incurs about 2 dB degradation [3].

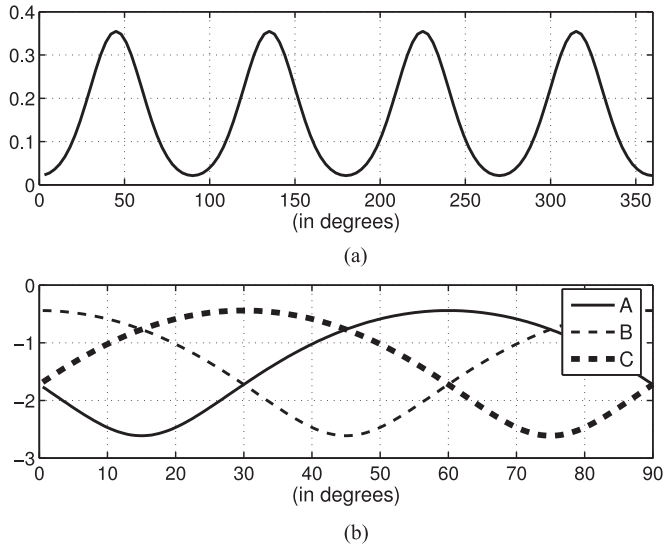

Fig. 2. (a) Probability Density of unquantized phase u at $\beta = 0$, $f_u(\alpha)$ for SNR $= 5$ dB. (b) Single step likelihoods $l(\phi|m)$ given $z = m$ and $\theta = 0^\circ$ ($M = 12$, SNR $= 5$ dB). A: $l(\phi|1) = l(\phi|4) = l(\phi|7) = l(\phi|10)$, B: $l(\phi|2) = l(\phi|5) = l(\phi|8) = l(\phi|11)$, C: $l(\phi|3) = l(\phi|6) = l(\phi|9) = l(\phi|12)$.

III. Phase Acquisition: Bayesian Estimation and the Need for Dither

Setting $\Delta f = 0$, the measurement model (1) specializes to

$$z_k = Q_M(u_k); \quad \text{where } u_k = \arg\left(b_k e^{j\phi} e^{-j\theta_k} + w_k\right) \quad (3)$$

where ϕ is the constant unknown channel phase offset. Since the noise is circularly symmetric, it is not affected by derotation. Given the model in (3), it is straightforward to derive the conditional pmf (probability mass function) of the observation, conditioned on the phase offset ϕ and derotation phase θ. It does not depend on k, and is denoted by $p(z = m|\phi, \theta) = p_\phi^\theta(z = m), m = \{1, ..., M\}$. The pmf is computed from (3) as follows. We first find the distribution of the unquantized phase u, $f_u(\alpha|\beta)$, conditioned on the value of the *net* rotation $\beta = \phi - \theta$. For a given QPSK symbol, u is simply the phase of a complex Gaussian random variable. The observation pmf is computed by integrating $f_u(\alpha|\beta)$ over appropriate bins. Details are provided in Appendix A.

Fig. 2(a) plots $f_u(\alpha|\beta = 0)$. We see that it is periodic with period 90° with modes at $\{45^\circ, 135^\circ, 225^\circ, 315^\circ\}$, because we choose the symbols uniformly from the QPSK constellation. It suffices, therefore, to limit ϕ to the interval $[0, 90^\circ)$. Fig. 2(b) shows the log likelihood plots $l(\phi|m) = \log\left(p_\phi^0(z = m)\right)$, as a function of the unknown phase ϕ, setting the derotation phase $\theta = 0$. Nonzero θ simply results in a circular shift, with

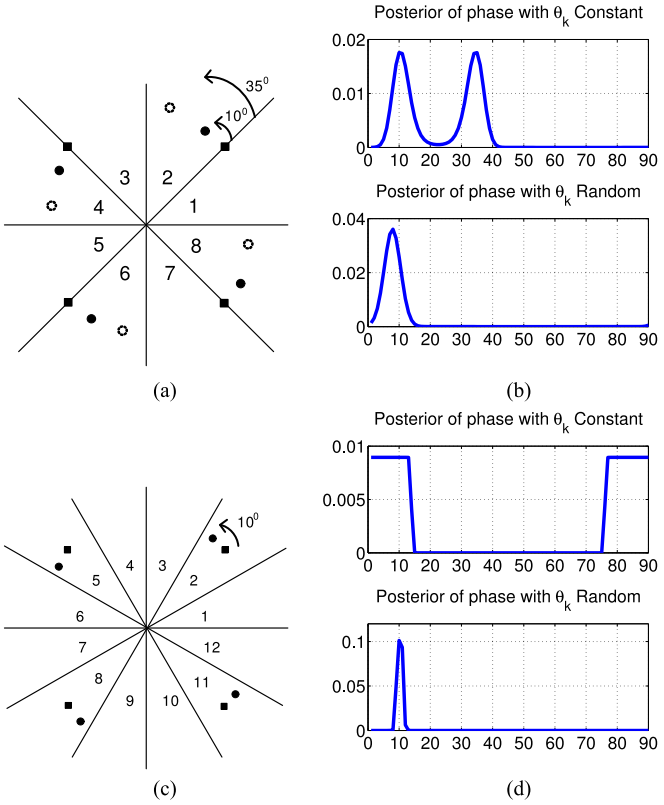

Fig. 3. (a) $M = 8$ uniform quantization regions. 4 1-bit ADCs at $\{I, Q, I \pm Q\}$. (b) Example 1: SNR $= 5$ dB and $M = 8$ regions. (c) $M = 12$ uniform quantization regions. 6 1-bit ADCs at $\left\{I, Q, I \pm \sqrt{3}Q, Q \pm \sqrt{3}I\right\}$. (d) Example 2: SNR $= 35$ dB and $M = 12$ regions. In subplots (a) and (c), solid black square dots denote the locations of transmitted QPSK symbols. Solid black round dots denote the noiseless symbol locations received after constant phase offset. (a) $\phi = 10°$, (b) Posterior for ϕ after 100 symbols (top) Derotation value θ_k kept constant (bottom) θ_k varied randomly, (c) $\phi = 10°$, (d) Posterior for ϕ after 30 symbols (top) Derotation value θ_k kept constant (bottom) θ_k varied randomly.

likelihood function given by $l(\phi - \theta|m)$, where it is understood that the argument is always expressed modulo 90°. An interesting property to note is the periodicity of $l(\phi|m)$ in the observation m, with period $M/4$. This follows from the symmetry induced by equiprobability of the transmitted QPSK symbols. For example, if $M = 8$ (Fig. 3(a)), only the observation modulo $M/4 = 2$ is relevant: $l(\phi|z) = l(\phi|\tilde{z})$, where $\tilde{z} = z \bmod 2 \in \{1, 2\}$.

A. Estimator Structure

Conditioned on the past derotation values θ_1^k (which are known) and the quantized phase observations z_1^k, applying Bayes rule and using independence of noise across symbols, we get a recursive equation for updating the posterior of the unknown phase as follows:

$$\begin{aligned} p\left(\phi|z_1^k, \theta_1^k\right) &\propto p\left(z_k|\phi, z_1^{k-1}, \theta_1^k\right) p\left(\phi|z_1^{k-1}, \theta_1^k\right) \\ &= p\left(z_k|\phi, \theta_k\right) p\left(\phi|z_1^{k-1}, \theta_1^{k-1}\right) \end{aligned} \quad (4)$$

Normalizing the probability density obviates the need to evaluate the denominator. Going to the log domain, we obtain an additive update for the cumulative log likelihood. Denoting by $l_{1:k}(\phi) = \log\left(p\left(\phi|z_1^k, \theta_1^k\right)\right)$ the cumulative update up to the kth symbol, we obtain a simple recursive update, as follows:

$$l_{1:k}(\phi) = l_{1:k-1}(\phi) + l(\phi - \theta_k) \quad (5)$$

The maximum a posteriori (MAP) estimate after N symbols is given by

$$\hat{\phi}_{\text{MAP};N} = \text{argmax} \qquad p\left(\phi|z_1^N, \theta_1^N\right) = \text{argmax} \qquad l_{1:N}(\phi)$$

We start with a uniform prior $p(\phi)$ over $[0°, 90°)$. Single step likelihoods, $l(\phi|m)$ for $m = 1, ..., M/4$, can be precomputed and stored offline. The recursive update (5) requires only the latest posterior to be stored.

B. The Need for Dither: Two Examples

As investigated in the next section, appropriate choice of the derotation phases provides a means of applying a *controlled dither* prior to quantization in order to aid in phase estimation. We motivate this in this section by considering two scenarios in which not applying dither (i.e., setting θ_k to a constant for all k) yields poor performance.

Example 1: Consider $M = 8$ phase quantization bins and $\phi = 10°$ (Fig. 3). In this case, not dithering ($\theta_k \equiv 0$) results in a spurious peak at $\phi = 35°$. We have already noted that, for $M = 8$, the observation z can be reduced to $\tilde{z} = z$ modulo 2 (i.e., noting whether we fall in an even or odd bin). Next, we note that circularly symmetric noise is equally likely to rotate us clockwise or anticlockwise. These two observations can be used to show that there is an unresolvable ambiguity in the likelihood function: $l(\phi|\tilde{z}) = l(45° - \phi|\tilde{z})$. For zero dither, this implies that the posteriors for ϕ and $45° - \phi$ are identical for any sequence of measurements. This ambiguity is formally described later in Lemma 2. Such ambiguities were also noted in the block noncoherent system in [3]. One approach to alleviate this ambiguity is to dither θ_k randomly; this dithers the spurious peak in the posterior while preserving the true peak, leading to a unimodal posterior distribution when computed over multiple symbols. Another approach is to break the symmetry in the phase quantizer, using 12 phase bins instead of 8. However, even this strategy can run into trouble at very high SNR, as shown in the next example.

Example 2: Now consider $M = 12$ phase bins and no noise (or very high SNR), again with true phase offset $\phi = 10°$. Since there is no noise, all observations fall in bins 2, 5, 8, 11, resulting in a flat phase posterior over the interval $[75°, 90°] \cup [0°, 15°]$ if there is no dither ($\theta_k \equiv 0°$) (for a formal statement see Lemma 1). This could lead to an error as high as 25° (Fig. 3). On the other hand, using randomly dithered θ_ks results in an accurate MAP estimate, with the combination of shifted versions (shifted by θ_k) of the flat posterior leading to a unimodal posterior with a sharp peak.

IV. Feedback Control for Phase Acquisition

While randomly dithered derotation is a robust design choice which overcomes the shortcomings of a naive no-dither strategy, it does not utilize the information obtained from the

measurements. It is natural to ask, therefore, whether we can do better with feedback control of the dither, with the goal of reducing the mean squared error (MSE) of the phase estimate faster. This problem of dither design falls in a general category of problems in *control for sequential estimation*, which has received significant attention recently in the context of multi-hypothesis testing [12]–[14], as well as for estimation of continuous-valued parameters [15]. Such problems are either solved over a finite horizon, in which case the goal is to minimize a metric such as the MSE, or over a variable horizon (with some stopping criterion), in which case the cost function to be minimized is the sum of the expected number of observations, plus a penalty term associated with the quality of the final estimate (e.g., the MSE). As discussed in the literature, either formulation can be mapped to a Partially Observable Markov Decision Problem (POMDP) whose optimal solution is computationally intractable. Significant recent effort [12]–[15] has therefore gone into characterizing asymptotically optimal solutions (in the limit of a large number of observations and a large coefficient for the penalty term). Since we are interested in phase estimation over a small number of observations, these results do not directly apply to our setting. However, the intuitively pleasing *Greedy Entropy Policy (GE)* policy that we employ is closely related to policies that have been used to derive theoretical bounds for multi-hypothesis testing [12].

Our GE policy picks an action at each step that minimizes the *expected* entropy (an information theoretic measure of uncertainty) of the next step phase posterior. It can be shown to be equivalent to a policy which, at each step, maximizes the mutual information between the new observation and the unknown phase offset. In this form, it is identical to a policy recently discussed in [14], [16] for hypothesis testing, where the goal is to maximize mutual information between the unknown hypothesis and the set of observations over a finite horizon. It is shown in [14], [16] that the greedy approach is the best among all polynomial time algorithms, and achieves a cost function which is within a constant factor $1/e$ of the optimal cost. While such guarantees translate to our problem as well, our interest is in minimizing MSE rather than maximizing mutual information.

A policy [15] that is closely related to ours is to maximize the Fisher information at each step, taking the latest MAP estimate as the true value of the parameter. While this Maximum Fisher Information (MFI) policy is shown to be asymptotically optimal in [15] under appropriate consistency conditions, its performance for a small number of observations is not investigated in [15]. We find that GE outperforms MFI in the latter regime, especially at low SNR, while converging to it (and hence inheriting its asymptotic optimality) as the number of observations gets large.

In this section, we first discuss the GE and MFI policies assuming consistency of the MAP estimate (i.e., assuming that, even with constant action, the posterior converges to a unimodal distribution centered around the true value of phase). This always holds for $M = 12$ with nonzero noise (Sub-Section IV-D). We then analyze the special case of zero noise separately, when the phase posteriors are flat and the MAP estimate is ill-defined. We show that in this case GE reduces the support of the posterior density by half at every step, thereby reducing the absolute error at an exponential rate. Finally, we discuss a simple strategy, mixing feedback control with intermittent random actions, for ensuring a consistent unimodal posterior when $M = 8$.

A. Greedy Entropy Policy

At step $k-1$ (i.e., after observing $k-1$ symbols) the net belief about the phase is captured by the posterior $f_{k-1}(\phi) := p\left(\phi | z_1^{k-1}, \theta_1^{k-1}\right)$. We now drop the subscript k, since the development below applies for any k. The entropy of the *current belief*, $f(\phi)$ is given by

$$h(f(\phi)) = -\int f(\phi) \log(f(\phi))\, d\phi \tag{6}$$

The new posterior, conditioned on the next action $\theta = \theta_k$ and observation $z = z_k$, is given by

$$f_{\text{new}}(\phi|\theta, z) = \frac{p_\phi^\theta(z)\, f(\phi)}{p^\theta(z)} \tag{7}$$

where $p_\phi^\theta(z)$ represents the conditional distribution of the observation (21) given the true phase offset, ϕ, and the derotation action, θ. The normalization term in the denominator is the probability density of observing z in the next step under the effect of taking action θ, averaged over the current belief, i.e.,

$$p^\theta(z) = \int p_\phi^\theta(z)\, f(\phi)\, d\phi \tag{8}$$

We can now compute the *expected* entropy of the new posterior if action θ is chosen, by averaging over the observation density $p^\theta(z)$ as follows:

$$h^\theta(f_{\text{new}}(\phi)) = E_z[h(f_{\text{new}}(\phi|\theta, z))] = \sum_{i=1}^{M} p^\theta(z_i) h(f_{\text{new}}(\phi|\theta, z)) \tag{9}$$

The GE policy chooses the derotation phase that minimizes the entropy of the new posterior, i.e.,

$$\theta_{GE} = \underset{\theta}{\text{argmin}} \quad h^\theta(f_{\text{new}}(\phi)) \tag{10}$$

This can also be expressed as maximization of *information utility*, IU^θ, which is the amount by which the uncertainty (entropy) is decreased due to the action θ :

$$\Rightarrow \theta_{GE} = \underset{\theta}{\text{argmax}} \left(h(f(\phi)) - h^\theta(f_{\text{new}}(\phi))\right) = \underset{\theta}{\text{argmax}}\, IU^\theta \tag{11}$$

This can in turn be expressed in terms of the Kullback-Leibler (KL) divergence (which is useful for proving the convergence of GE to MFI as discussed later), using (11), (6), (7):

$$IU^\theta = \int f(\phi)\, D^\theta(\phi)\, d\phi \tag{12}$$

where $D^\theta(\phi)$ is the KL divergence between densities $p_\phi^\theta(z)$ and $p^\theta(z)$

$$D^\theta(\phi) = \sum_i p_\phi^\theta(z_i) \log \frac{p_\phi^\theta(z_i)}{p^\theta(z_i)} \tag{13}$$

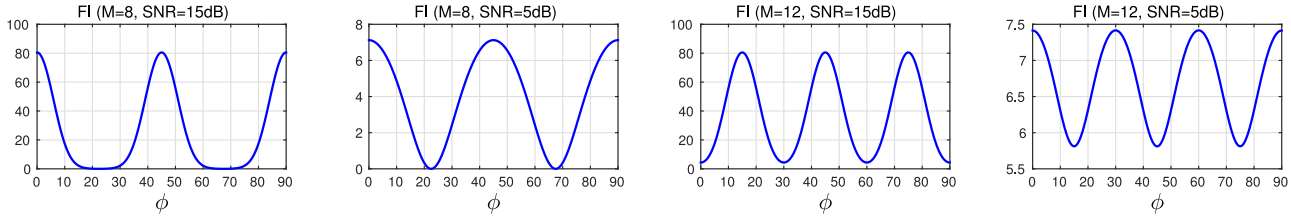

Fig. 4. Fisher Information as a function of ϕ ($\theta = 0$).

It is straightforward to implement the greedy entropy policy by evaluating the information utility (12) over the finite set of actions (i.e., the discretized set of phases from which the dither is chosen).

B. Fisher Information

Fisher information provides a measure of the *sensitivity* of the estimation problem to the value of the parameter being estimated; parameter values that result in higher Fisher information can be estimated with greater accuracy or with fewer measurements. The well-known Cramer-Rao bound, which is the inverse of the Fisher information, provides a lower bound on the mean square error for any unbiased estimator. For our phase estimation problem, the Fisher information as a function of the true phase ϕ and the derotation action θ, is given by:

$$FI^{\theta}(\phi) = \sum_{i=1}^{M} \left(\frac{\partial p_{\phi}^{\theta}(z_i)}{\partial \phi} \right)^2 \cdot \frac{1}{p_{\phi}^{\theta}(z_i)} \tag{14}$$

The derivative of the observation density $p_{\phi}^{\theta}(z)$ can be computed by differentiating the function $f_u(\cdot)$ prior to integration; see (20) and (21) in Appendix A. Given the expression in (20), we note that evaluating the derivative of $f_u(\cdot)$ wrto ϕ is straightforward as it comprises of easily differentiable functions (including *erfc*). In Fig. 4 we plot the Fisher information as a function of the phase offset (setting the dither $\theta = 0$) for 4 different cases: SNR low or high and number of quantization bins $M = 8, 12$. We observe that in three of the cases, Fisher information is maximum for phase offsets that bring the final phase after rotation to the "boundary" i.e., one of the bin edges. This is intuitive at high SNR: if the complex QPSK symbol ends up being in the "middle" of the quantization bin, the same measurement would be recorded at every symbol period, resulting in a flat posterior and hence a poor estimate. Interestingly, when the noise is high enough to knock the symbol around more and the bins are narrower ($M = 12$), Fisher information is maximized for a phase offset (30°) that brings the symbol to the "middle" of the quantization cone (Fig. 4(d)) (for instance, if the QPSK symbol $\frac{\pi}{4}$ is transmitted, the net phase is $30^\circ + 45^\circ = 75^\circ$ which is exactly in between the phase thresholds at angles 60° and 90°).

Genie optimal lower bound: The preceding Fisher information computations provide us with a "genie" optimal control policy; the best action for any ϕ is the one that brings the net phase $\phi - \theta$ to a value for which the Fisher information is maximized. This yields a Cramer-Rao bound which provides a lower bound for MSE against which any policy can be compared.

Since we do not know the actual value of the phase ϕ, a natural approach is to use the best guess, which is the latest MAP estimate. This is the MFI policy, which chooses actions at each step as follows:

$$\theta_{\mathrm{MFI}} = \underset{\theta}{\operatorname{argmax}}\ FI^{\theta}(\phi_{\mathrm{MAP}})$$

where

$$\phi_{\mathrm{MAP}} = \underset{\phi}{\operatorname{argmax}}\ f(\phi) \tag{15}$$

where $FI^{\theta}(\phi)$ is computed via (14). $f(\phi)$ is the latest belief/posterior distribution of the phase offset. MFI chooses near-optimal actions if the MAP estimate is close to the true offset, and is therefore asymptotically optimal under consistency assumptions [15]. However, when the uncertainty in $f(\phi)$ is high (and the MAP estimate is poor), we expect a policy that takes into account the entirety of the posterior distribution (rather than just its maximum), such as the GE, to perform better. This is borne out by simulation results presented shortly. On the other hand, as the uncertainty in $f(\phi)$ reduces and the estimator becomes more confident of the MAP estimate, the GE policy reduces to MFI under a Gaussian approximation for the posterior, as stated in the following theorem.

Theorem 1: Suppose that the latest phase posterior is normally distributed, i.e., $f(\phi) \sim \mathcal{N}(\phi_0, v^2)$ where v is in the unit of radians. Then, as $v \to 0$, the GE policy chooses the same actions as the MFI policy, i.e.,

$$\lim_{v \to 0} \ \underset{\theta}{\operatorname{argmax}} \ IU^{\theta} = \underset{\theta}{\operatorname{argmax}} \ FI^{\theta}(\phi_0) \tag{16}$$

Specifically

$$\lim_{v \to 0} \ \frac{IU^{\theta}}{v^2} = \frac{1}{2} FI^{\theta}(\phi_0) \tag{17}$$

The proof is provided in Appendix (B). Of course, $f(\phi)$ is not strictly Gaussian as its support is $\left[0, \frac{\pi}{2}\right]$, but under consistency and asymptotic normality, the results kick in as the number of observations increases. In fact, in our simulations, we find that the equation $\operatorname{argmax}_{\theta}\ IU^{\theta} \approx \operatorname{argmax}_{\theta}\ FI^{\theta}(\phi_0)$ starts holding as soon as the standard deviation of $f(\phi)$ is within a few degrees. We also note from the theorem that the value of the information utility scales with the variance of the posterior density, independent of the actions.

C. The Zero Noise Setting

As discussed earlier, when SNR is very high, the resulting posterior density is flat over a support interval determined by

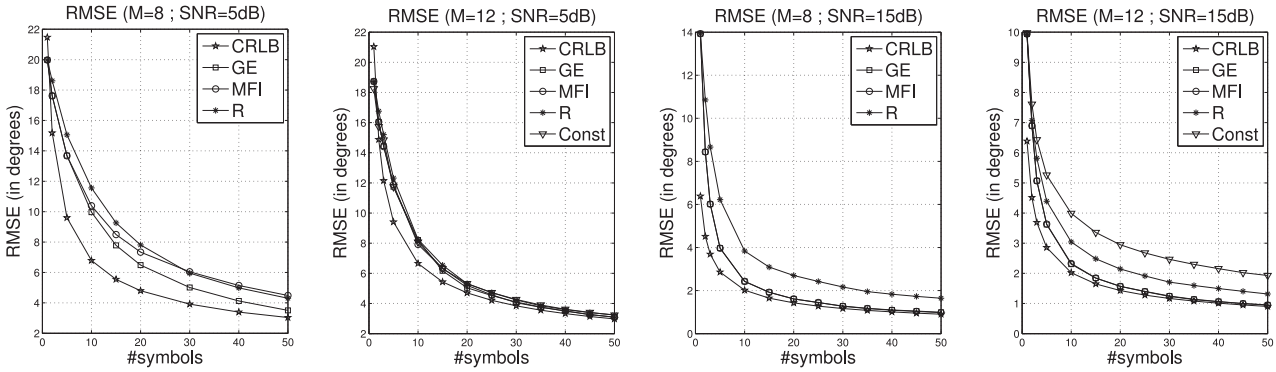

Fig. 5. Results of Monte Carlo simulations of different strategies for choosing the feedback θ_k with 4 and 6 ADCs (8 and 12 phase bins) at SNRs 5 dB and 15 dB. Policies: Greedy Entropy (GE), Maximizing Fisher Information (MFI), Random dither (R) and Constant derotation phase (Const).

the set of observations. In this case, dithering is critical, since fixed θ_k results in the same measurement (modulo the symmetries induced by the constellation) and no change in posterior. This is a common feature in systems involving heavily quantized measurements: at high SNR, dither acts as artificial noise and provides the necessary diversity of measurements required for estimation. In this zero noise setting, the posterior remains always flat, but its support changes as we change the action. GE is equivalent to choosing the action that reduces the support the most, and is therefore optimal. This is established via the following lemma, whose proof is sketched in Appendix (C).

Lemma 1: In the absence of noise (i.e., $w_k = 0 \quad \forall k$ in (3)), the phase posterior $f_k(\phi)$ is uniform over its support for each k. Let S_k denote the size of its support at time k. The action chosen by the Greedy Entropy policy is the one that minimizes the expected value of S_{k+1}. Furthermore, $S_{k+1} = \frac{1}{2}S_k$, hence the absolute phase error reduces exponentially at the rate of $\frac{1}{2}$. MFI is not well defined as there is no unique MAP estimate, but if the MMSE estimate is used instead in (15), then MFI chooses the same actions as the GE policy.

D. Avoiding Phase Ambiguity for $M = 8$

We have assumed thus far that the MAP estimate converges to the correct phase offset irrespective of the sequence of actions taken. This is indeed true for $M = 12$. This is because for any action θ, different values of the true phase offsets result in distinct observation densities. This is expressed mathematically in terms of KL Divergence as follows

$$\text{for any } \phi \neq \phi', \quad D\left(p_{\phi}^{\theta} || p_{\phi'}^{\theta}\right) > 0 \quad \forall\, \theta\, (M = 12) \qquad (18)$$

However, the preceding condition does not hold for $M = 8$. Due to the symmetry of the angular thresholds, for any given value of ϕ and a given derotation θ, there exists another phase offset, ϕ', which results in an identical distribution over the quantized measurements. This means that if θ is kept constant, the limiting posterior $f(\phi)$ is bimodal, with true and spurious peaks at locations ϕ and ϕ' respectively. The value of ϕ' is a function of ϕ (which remains fixed) and θ. The lemma below specifies this relationship.

Lemma 2: When $M = 8$ and the true phase is denoted by $\phi \in \left[0, \frac{\pi}{2}\right)$, for any derotation phase θ, there exists a value $\phi' \in \left[0, \frac{\pi}{2}\right) \neq \phi$, such that $D(p_{\phi}^{\theta} || p_{\phi'}^{\theta}) = 0$. This holds for $\phi' = \text{mod}\left(2\theta - \phi + \frac{\pi}{4}, \frac{\pi}{2}\right)$.

The proof is provided in Appendix (D). We see that a constant dither policy is unacceptable as it leaves a bimodal ambiguity in the value of the phase offset. A random dither continuously changes θ and thereby *guarantees* a unimodal limiting posterior. However, feedback control policies such as GE or MFI, which typically also eliminate bimodality, may occasionally run into trouble, with a small probability of the final posterior being maximized at the *spurious* phase offset value. This can happen in the following manner: suppose a total of N measurements are made, out of which a majority, say $N_1 \approx N$ employed a constant action (this can happen, say with MFI if ϕ_{MAP} remains same). In the remaining few steps, $N_2 = N - N_1$, different value(s) of θ were used. Recall that the final ϕ posterior is just a summation of the individual step log likelihoods, the order being irrelevant. Now it may happen that these few N_2 observations are affected by bad noise instances and the ϕ posterior, computed based on just these steps, has a larger probability mass at the spurious value. Since the posterior distribution from the other N_1 steps is perfectly bimodal, the combined posterior ends up having a stronger peak at ϕ'. The probability of such an event is generally very small, as it requires getting multiple bad measurements during which ϕ' should appear to be more probable. However, it does occur occasionally during our Monte Carlo runs.

Fortunately, a simple modification to the GE and MFI policies can guarantee vanishing probabilities for such bad events. The idea is to pick the actions randomly for a fixed fraction, γ, of the steps. For simplicity, consider inserting such actions at regular intervals; for instance, $\gamma = 0.1$ means choosing every 10th action randomly, while the remainder are chosen in the usual manner as dictated by the feedback control policy being employed. As N gets large, so does the number of random dither steps, γN (for $\gamma > 0$), thereby ensuring that the limiting posterior is unimodal and that the MAP estimate converges to the correct phase. Note that a more efficient scheme can also be used, as described in the [13], where the randomly chosen actions are scheduled at intervals which grow exponentially. However, for the small number of measurements (typically less

than 100) of interest to us, the fixed rate schedule for inserting random actions works well, with no noticeable deterioration in the efficiency of the feedback control policy.

E. Simulation Results

The root mean squared error (RMSE) performance of phase acquisition is evaluated using Monte Carlo simulations averaging over randomly generated channel phases. Fig. 5 plots results for two values of SNR: a low value of 5 dB and a high value of 15 dB. Errors are computed modulo 90°, for instance if the true phase offset is 80° and the estimate is 5°, this is equivalent to an error of 15°. We implement three policies: greedy entropy (GE), random dither (R) and maximizing the Fisher information (MFI). We also simulate the policy of keeping the derotation phase constant when $M = 12$, the case for which it is consistent. For comparison we plot the genie-optimal performance, which is the CRLB computed by inverting the maximal Fisher information (maximized over the true phase offset ϕ, setting $\theta = 0$). However, note that this does not give a valid lower bound when the number of measurements are few and the errors can be large. This is because the Cramer-Rao bound is based on the standard notion of squared error, not the modulo 90° error appropriate in our setting. In principle, such problems could be addressed via a tighter and more sophisticated bound. However, even for a moderate number of observations, we find that the error reduces quickly enough that it becomes unnecessary to distinguish between the two notions of computing error.

From the plots, we make the following observations: (a) The performance of GE is very close to the "genie" optimal control policy (CRLB) in all cases. (b) The performance of GE and MFI is almost identical, but GE is slightly better at low SNR and coarser quantization (5 dB, 8 bins), when the MAP estimate that MFI relies upon can be poor initially. (c) At low SNR, there is little to distinguish between random dithering and GE, since the noise supplies enough dither to give a rich spread of measurements across different bins. In fact at low SNR and finer quantization (5 dB, 12 bins), constant action performs as well as others. However, when the quantization is more severe (8 bins), the GE policy provides performance gains over random dithering even at low SNR. To summarize, we find that efficient dithering policies could be effective for rapid phase acquisition under the scenarios of more severe quantization and higher SNRs.

Once an accurate enough phase estimate is obtained in the acquisition step, we wish to begin demodulating the data, while maintaining estimates of the phase and frequency. In the next section, we describe an algorithm for decision directed (DD) tracking. In this DD mode, the phase derotation values θ_k aims to correct for the net channel phase in order to enable accurate demodulation. This is in contrast to the acquisition phase discussed so far, where the derotation is designed to aid in phase estimation.

V. Phase/Frequency Tracking

We must now account for the frequency offset in order to track the time-varying phase, and to compensate for it via derotation in order to enable coherent demodulation. As we have discussed, the phase is well approximated as roughly constant over a few tens of symbols, whereas accurate estimates of the frequency offset η (2) require observations spanning hundreds of symbols. This motivates a hierarchical tracking algorithm. Bayesian estimates of the phase are computed over relatively small windows, modeling it as constant but unknown. The posterior computations are as in the acquisition stage, with two key differences: the derotation phase value is our current best estimate of the phase, and we operate in decision-directed mode, and hence do not need to average over all possible symbols. These relatively coarse phase estimates are then fed to an extended Kalman filter (EKF) for tracking both frequency and phase. The filter is initialized with the phase estimate from the acquisition stage. The data is differentially encoded over the QPSK symbols (this is necessary, since phase estimation was performed modulo $\frac{\pi}{2}$ in the acquisition stage).

Denote by $\hat{\phi}_{\mathrm{MAP};W}(k)$ the MAP phase estimate over a sliding window of W symbols. This is fed as a noisy measurement of the true time varying phase $\phi(k)$ to an EKF constructed as follows:

Process Model:

$$x_k = Ax_{k-1} + w_k$$

$$\begin{bmatrix} \phi(k) \\ \eta(k) \end{bmatrix} = \begin{bmatrix} 1 & T_s \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \phi(k-1) \\ \eta(k-1) \end{bmatrix} + w(k)$$

where $w(k) \sim \mathcal{N}(0, Q_{\mathrm{LO}})$. We set the noise covariance matrix Q_{LO} using the two-state model for the LO (local oscillator) clock dynamics, as discussed in [17].

$$Q_{\mathrm{LO}} = w_c^2 q_1^2 \begin{bmatrix} T_s & 0 \\ 0 & 0 \end{bmatrix} + w_c^2 q_2^2 \begin{bmatrix} \frac{T_s^3}{3} & \frac{T_s^2}{2} \\ \frac{T_s^2}{2} & T_s \end{bmatrix} \tag{19}$$

$w_c = 2\pi f_c$ represents the carrier frequency (in rad/s) and parameters q_1^2 (units of seconds) and q_2^2 (units of Hertz) are the noise parameters corresponding to white frequency noise and random walk frequency noise respectively. As discussed in [17] their values can be determined from the Allan variance of the LO, which in turn can be computed from the LO phase noise characteristics [18].

Measurement Model:

$$y_k = h(x_k) + v_k$$

$$y(k) = \begin{bmatrix} \cos\left(4 \cdot \hat{\phi}_{\mathrm{MAP};W}(k)\right) \\ \sin\left(4 \cdot \hat{\phi}_{\mathrm{MAP};W}(k)\right) \end{bmatrix} = \begin{bmatrix} \cos(4 \cdot \phi(k)) \\ \sin(4 \cdot \phi(k)) \end{bmatrix} + v(k)$$

where $h(\cdot)$ is a nonlinear measurement function. The particular form is chosen to avoid explicit phase unwrapping. Since tracking is done in decision directed mode, there is no need to average over the distribution of QPSK symbols (this also removes the ambiguity that was present with $M = 8$ during acquisition) and thus the phase $(\phi(k))$ is estimated over the interval $[0, 2\pi)$. However as differential encoding is being used, integer shifts of 90° in the phase estimate are permissible, hence a factor of 4 is used inside the sine and cosine arguments. The measurement noise is $v(k) \sim \mathcal{N}(0, R_k)$. For the EKF, computation of the

Jacobian of the nonlinear function $h(\cdot)$ is required, which in this case evaluates to

$$H_k = \begin{bmatrix} -4\sin(4\phi(k)) & 0 \\ 4\cos(4\phi(k)) & 0 \end{bmatrix}$$

The EKF update equations are given as follows (these are standard EKF equations, we refer the readers to Chapter 10 of [19] for a discussion on EKF).

Time Update:

$$\hat{x}_{k|k-1} = A\hat{x}_{k-1}; \quad \hat{P}_{k|k-1} = A\hat{P}_{k-1}A^T + Q_k$$

$$K = \hat{P}_{k|k-1}H_k^T\left(H_k\hat{P}_{k|k-1}H_k^T + R_k\right)^{-1}$$

Measurement Update:

$$\hat{x}_k = \hat{x}_{k|k-1} + K(y_k - h(\hat{x}_{k|k-1}))$$

$$\hat{P}_k = (I - KH_k)\hat{P}_{k|k-1}$$

where $\hat{P}_k$ is the estimate of the state error covariance and H_k is evaluated at $\hat{x}_{k|k-1}$. The *cleaned* state estimate, $\hat{x}_k$, provides the *latest* estimate of the frequency offset $\hat{\eta}(k) = \hat{x}_k(2)$ and a *delayed* estimate of the net phase, delayed due to the effect of sliding window. The measurement at time k, y_k, reflects the phase estimated over the time window $[k-W, k]$, hence the feedback (for undoing the phase at time k) is set according to $\theta_k = \hat{x}_k(1) + \frac{W}{2}T_s\hat{\eta}(k)$.

A. Setting the Noise Covariances

To compute a practically relevant value of Q_{LO} for simulations, we use the specifications of a Hittite 60 GHz receiver [20]. The phase noise characteristics in the specification sheet are used to compute the Allan variance, which gives the process noise parameter values $q_1^2 = 1.23 \cdot 10^{-22}$ s and $q_2^2 = 3.5 \cdot 10^{-21}$ s^{-1}. This leads to a noise covariance matrix with very small variances ($Q_{\mathrm{LO}} \approx [3 \cdot 10^{-9}$ rad^2, 0; 0, 10^{-7} (rad/s)^2]). Thus, the process noise for typical oscillators is small and can be tracked very easily. In order to enable the filter to react to abrupt changes in the value of frequency offset (e.g., due to changes in Doppler), we *artificially inflate the process noise.* We can afford to do so because the resulting marginal increase in phase offset error has a negligible effect on BER, which is our ultimate performance measure. For the measurement noise, the covariance matrix R_k depends on the uncertainty, σ_ϕ^2, in the *windowed* MAP phase estimate. While the latter can be estimated empirically from the posterior, we find it convenient to use an analytical approximation that is in close agreement with empirical estimates. For the approximation, we decompose the phase error into two independent contributions, and set $\sigma_\phi^2 = \sigma_{\text{fixed}}^2 + \sigma_{\text{sliding}}^2$. The first term corresponds to estimation error, assuming that the true value remains fixed over the estimation window, and is well approximated by the CRLB for classical phase estimation [21]: $\sigma_{\text{fixed}}^2 = \frac{\sigma^2}{W}$ (i.e., this first terms depends only on SNR and estimation window size W). The second term represents the error in the piecewise constant phase model due to the linear phase change uncertainty due to the frequency offset, and evaluates to $\sigma_{\text{sliding}}^2 = \frac{\eta^2 T_s^2 (W^2-1)}{12}$, where

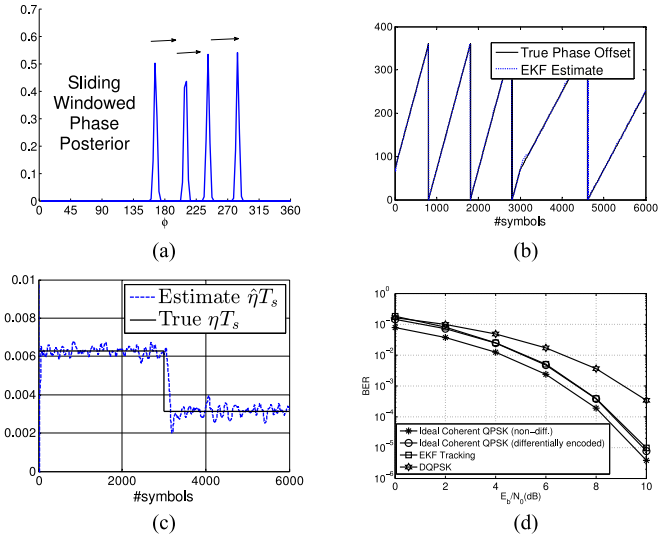

Fig. 6. Performance plots of EKF based Tracking Algorithm. (a) posterior of phase (SNR = 6 dB), (b) Phase Estimates (SNR = 6 dB), (c) Frequency offset estimate ($\hat{\eta}T_s$) (at SNR = 6 dB), (d) Bit Error Rate Plots.

value of frequency offset η is plugged in from the EKF estimate. Finally, in order to compute the entries of the measurement noise covariance R_k (i.e., Var$[\cos(4\phi(k))]$, Var$[\sin(4\phi(k))]$ and Cov$[\cos(4\phi(k))\sin(4\phi(k))]$), we make the simplifying assumption that $\phi(k)$ is normally distributed with mean $\hat{\phi}_{\mathrm{MAP};W}(k)$ and variance σ_ϕ^2. This Gaussianity assumption is a good fit to the empirical posterior distribution (Fig. 6(a)), and enables straightforward analytical computation of the entries of R_k (expressions omitted due to lack of space). The analytical estimates obtained are close to simulation results (when the variance estimate from the posterior distribution is used directly).

B. Simulation Results

We use $M = 8$ bins and sliding window length $W = 40$. The EKF algorithm accurately tracks the phase (6(b)). Subplot 6(a) shows several superimposed snapshots of the windowed posterior of the phase, whose peaks (the MAP estimates) are used as measurements for the EKF. In subplot 6(c) ηT_s was changed from $2\pi \cdot 10^{-3}$ rad/symbol to $\pi \cdot 10^{-3}$ rad/symbol after 3000 symbols ($T_s = \left(6 \times 10^9\right)^{-1}$ secs). The plot shows $\hat{\eta}T_s$, the estimate, by setting Q_{LO} to be $[3 \cdot 10^{-9}$ rad^2, 0; 0, $2 \cdot 10^9$ (rad/s)^2] (with noise inflated in the frequency offset term), which enables the filter to lock onto the new value in about 200 symbols. Subplot 6(d) shows BER curves for ideal (unquantized) differentially encoded coherent QPSK and that of the proposed algorithm, which is very close to the former. 400 runs with 25000 long bit sequences each were used to generate the BER plots. Using noncoherent differential QPSK (DQPSK) obviates the need for phase synchronization but results in a 2 dB performance degradation.

VI. Conclusion

Hybrid analog-digital architectures with feedback provide a promising approach for DSP-centric designs that exploit

Moore's law, by alleviating the ADC bottleneck encountered at high communication bandwidths. In this paper, we show that a simple digitally controlled analog preprocessing step prior to quantization enables efficient use of the limited number of ADCs available for phase quantization. The derotation feedback provides (a) the *dither* required for phase acquisition with coarse quantization and high SNR and (b) the *correction* required to keep the received symbols in the center of the decision boundaries for optimal coherent demodulation in the tracking step.

The Bayesian framework as described in the paper effectively handles the joint uncertainties of data and the channel for a system with significant nonlinearities due to quantization. The posteriors are useful for computing both the feedback for the analog preprocessor and the quantities to be estimated. The BER obtained by our approach is comparable to that of an unquantized QPSK with differential decoding, unlike the degradation of performance in an open loop system such as [3].

While the problem of optimal dither for blind acquisition is a POMDP which is computationally intractable in general, we show that the intuitively appealing strategy of choosing derotations for minimizing the entropy of phase posterior is close to genie-optimal over the short time windows of interest to us, and prove that it is asymptotically optimal over large time windows.

An important direction for future work is to explore the challenges of implementing such ideas in more specific settings, as well as to explore the fundamentals of mixed-signal strategies for alleviating the ADC bottleneck in more challenging settings. The phase-only quantization strategy considered here suffices for PSK constellations over channels with minimal dispersion, but more complex approaches are required to handle channel dispersion and automatic gain control (the latter is important for amplitude-phase constellations).

In this paper we have presented results assuming ideal sampling. Simulations (skipped in the paper in favor of ease of exposition) show that our approach is robust for small amounts of timing mismatches ($<10\%$). Larger timing offsets result in large ISI (inter symbol interference) which needs to be handled separately. This is part of our future work.

Appendix A
Observation Density

The unquantized phase is given by $u = \arg\left(e^{j(2i-1)\frac{\pi}{4}}e^{j\beta} + w\right)$, where i is uniformly distributed over $\{1,2,3,4\}$ and $\beta = \phi - \theta$. It is straightforward to evaluate the density of the argument of a Gaussian random variable [10], we get the following expression for the density of u:

$$f_u(\alpha;\beta) = \sum_{i=1}^{4} \frac{1}{4}\left[\frac{a_i\left(2-\operatorname{erfc}\left(\frac{a_i}{\sigma\sqrt{2}}\right)\right)e^{\frac{a_i^2-1}{2\sigma^2}}}{2\sigma\sqrt{2\pi}} + \frac{e^{-\frac{1}{2\sigma^2}}}{2\pi}\right]$$

where

$$a_i = \cos\left((2i-1)\frac{\pi}{4} + \beta - \alpha\right) \tag{20}$$

Given the density of u, the observation pmf can be evaluated as follows:

$$p_\phi^\theta(z=m) = P(z=m|\beta) = \int_{(m-1)\frac{2\pi}{M}}^{m\frac{2\pi}{M}} f_u(\alpha;\beta)\ d$$

where

$$m \in \{1,2,.....,M\} \tag{21}$$

The pmf for $\beta = 0^\circ$ is evaluated numerically and stored. Since the dependence on β is only through the cosine, pmf at nonzero β values can be obtained by simple circular shifts of $P(z|\beta = 0^\circ)$.

Appendix B
Proof of Theorem 1

The equivalence of Fisher information to the second derivative of Kullback-Leibler divergence between two parametric densities with small perturbations is well known [22]. In this proof we encounter a similar relation. Consider the Taylor series expansion of the KL divergence (13) centered at ϕ_0 (note that $\phi_0 = \phi_{\mathrm{MAP}}$ since $f(\phi) \sim \mathcal{N}(\phi_0, v^2)$)

$$D^\theta(\phi) = D^\theta(\phi_0) + (\phi-\phi_0)D'^\theta(\phi_0) + \frac{(\phi-\phi_0)^2}{2}D''^\theta(\phi_0) + \cdots \tag{22}$$

the $\mathrm{superscripts}'$ and $apos;''$ denote derivatives with respect to ϕ. Substituting this in (12) gives

$$IU^\theta = D^\theta(\phi_0)\int f(\phi)\,d\phi + D'^\theta(\phi_0)\int f(\phi)(\phi-\phi_0)\,d\phi + D''^\theta(\phi_0)\int f(\phi)\frac{(\phi-\phi_0)^2}{2}d\phi + \cdots \tag{23}$$

since $f(\phi)$ is normally distributed, this simplifies to

$$IU^\theta = D^\theta(\phi_0) + \frac{v^2}{2}D''^\theta(\phi_0) + O(v^4) \tag{24}$$

$$\Rightarrow \lim_{v\to 0}\frac{IU^\theta}{v^2} = \lim_{v\to 0}\frac{D^\theta(\phi_0)}{v^2} + \lim_{v\to 0}\frac{1}{2}D''^\theta(\phi_0) \tag{25}$$

Consider the first term in the equation above

$$\frac{D^\theta(\phi_0)}{v^2} = \sum_i \frac{p_{\phi_0}^\theta(z_i)}{v^2}\log\left(\frac{p_{\phi_0}^\theta(z_i)}{\int p_\phi^\theta(z_i)f(\phi)\,d\phi}\right) = \sum_i \frac{p_{\phi_0}^\theta(z_i)}{v^2}\log\left(\frac{p_{\phi_0}^\theta(z_i)}{p_{\phi_0}^\theta(z_i) + \frac{v^2}{2}h_{\phi_0}^\theta(z_i) + O(v^4)}\right)$$

where

$$h_\phi^\theta(z) = \frac{\partial^2 p_\phi^\theta(z)}{\partial\phi^2} \tag{26}$$

where we have used the Taylor series expansion for $p_\phi^\theta(z_i)$ around ϕ_0 in (26). Applying the limit $v\to 0$ using the L'Hospital's rule (and using the fact that $p_\phi^\theta(z)$ is strictly positive

for any finite SNR), the expression above simplifies to

$$\lim_{v\to 0}\frac{D^{\theta}(\phi_0)}{v^2} = \frac{-1}{2}\sum_i h^{\theta}_{\phi_0}(z_i)$$
$$= \frac{-1}{2}\sum_i \frac{\partial^2 p^{\theta}_{\phi_0}(z_i)}{\partial\phi^2} = \frac{-1}{2}\frac{\partial^2}{\partial\phi^2}\left(\sum_i p^{\theta}_{\phi_0}(z_i)\right)$$
$$= \frac{-1}{2}\frac{\partial^2}{\partial\phi^2}(1) = 0$$

where we use the fact that $p^{\theta}_{\phi}(z)$ is the observation density and hence sums to 1. The first term in (25) is thus 0. For the second term, evaluating the double derivative of the KL divergence and using simple arithmetic simplifications (that we skip) gives

$$\frac{1}{2}D''^{\theta}(\phi_0) = \frac{1}{2}\sum_i h^{\theta}_{\phi_0}(z_i)\log\left(\frac{p^{\theta}_{\phi_0}(z_i)}{\int p^{\theta}_{\phi}(z_i) f(\phi)\, d\phi}\right)$$
$$+\frac{1}{2}\sum_i \left(\frac{\partial p^{\theta}_{\phi}(z_i)}{\partial\phi}\right)^2_{\phi=\phi_0}\frac{1}{p^{\theta}_{\phi_0}(z_i)}$$

which is a summation of two terms, the second one is the fisher information evaluated at ϕ_0.

$$\frac{1}{2}D''^{\theta}(\phi_0) = \frac{1}{2}T_1 + \frac{1}{2}FI^{\theta}(\phi_0) \tag{27}$$

Fisher information is independent of v. The proof of the theorem is complete by observing that the first terms goes to 0 as $v\to 0$. This is because the argument of the log term approaches 1.

$$\lim_{v\to 0}\frac{p^{\theta}_{\phi_0}(z_i)}{\int p^{\theta}_{\phi}(z_i) f(\phi)\, d\phi} = 1 \tag{28}$$

This can be easily derived by using the Taylor series expansion of $p^{\theta}_{\phi}(z_i)$ around ϕ_0.

Appendix C
Proof of Lemma 1

The first part of the lemma follows directly from the recursive Bayes equation (4) and by noting that in absence of noise, the single step phase density, $p^{\theta}_{\phi}(z)$, is uniformly distributed over ϕ (with support $\frac{2\pi}{M}$) for any given value of θ and z.

Since $f_k(\phi) = \frac{1}{S_k}$ over its support and zero otherwise, its entropy is given by

$$h(k) = -\int f_k(\phi)\log(f_k(\phi))\, d\phi = \log(S_k) \tag{29}$$

i.e., the entropy of a uniform density is equal to the logarithm of the length of the support interval. Hence minimizing entropy corresponds to minimizing the support. Let us denote the support interval of $f_k(\phi)$ by $\left[\phi^1_k, \phi^2_k\right]$; $0 \le \phi^1_k \le \phi^2_k$ (we can assume it to be of this particular form if we do not wrap around to force the phase to lie in the interval $\left[0, \frac{\pi}{2}\right)$, something that we do in practice for a simpler implementation). Note that $\phi^2_k - \phi^1_k = S_k$ and $S_k \le \frac{2\pi}{M}$. Now, conditioned on the action θ_{k+1} and the QPSK symbol $p_k\frac{\pi}{4}$; $p_k \in \{1,3,5,7\}$, the net final phase in the next step, Ω_{k+1}, lies uniformly in the interval $\Omega_{k+1} \in \left[\Omega^1_{k+1}, \Omega^2_{k+1}\right] = \left[\phi^1_k - \theta_{k+1} + p_k\frac{\pi}{4}, \phi^2_k - \theta_{k+1} + p_k\frac{\pi}{4}\right]$. Since this interval is less than $\frac{2\pi}{M}$, the bin size, there are only two quantized phase measurements possible at $k+1$; let us denote them by indices $i-1$ and i (Fig. 7).

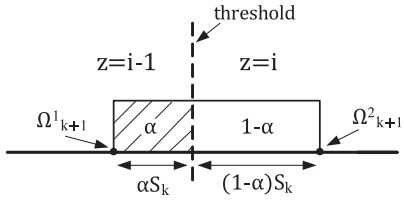

Fig. 7. Distribution of the net phase Ω_{k+1}. Dotted line denotes the phase threshold. Note that $\Omega^2_{k+1} - \Omega^1_{k+1} = S_k$.

$$p^{\theta_k}_{\phi}(z_{k+1}) = \begin{cases}\alpha, z_{k+1} = i-1\\ 1-\alpha, z_{k+1} = i\end{cases}$$
$$\alpha = Pr(\Omega_{k+1} \le \text{threshold}) \in [0,1]$$

The relative probabilities of getting these two measurements, denoted by $\{\alpha, 1-\alpha\}$, is determined by the action θ_{k+1} through which we can control the location of the uniform Ω_{k+1} density relative to the closest threshold. It can be easily seen that if we get the measurement $z_{k+1} = i-1$, the uncertainty in phase will be reduced to an interval of size αS_k. This means that the conditional entropy $h(k+1|z=i-1) = \log(\alpha S_k)$. Similarly $h(k+1|z=i) = \log((1-\alpha)S_k)$. Hence the average entropy is given by

$$E[h(k+1)] = \alpha\log(\alpha S_k) + (1-\alpha)\log((1-\alpha)S_k) \tag{30}$$

which is minimized for $\alpha = \frac{1}{2}$. This is achieved by the GE policy by choosing an action θ that places the net phase distribution symmetrically around one of the thresholds. Irrespective of the measurement, the support of the new posterior is half of the earlier support, i.e., $S_{k+1} = \frac{S_k}{2}$. Note that this strategy is optimal as choosing any value of α other than $\frac{1}{2}$ results in a support size that on average is greater than half of the previous support. Also note that even though MFI is not well defined because of the flat posterior, if we instead choose ϕ_{MMSE}, the mean of the posterior, it is equivalent to GE since fisher information is maximized when the net phase is placed at the "boundary" at high SNR.

Appendix D
Proof of Lemma 2

The key observation to see why the lemma holds is this: it can be easily inferred from (20) and (21) that the set of phase offset rotations $\beta = \phi - \theta = \left\{\alpha, \frac{\pi}{4} - \alpha + k\frac{\pi}{2}\right\}$; $k \in \mathbb{I}$; $\forall\ \alpha$ result in identical conditional densities $P(z|\beta)$ when M=8. For fixed derotation, these different values correspond to different phase offsets. Setting $k=0$ we can write:

$$\alpha = \phi - \theta \text{ and } \frac{\pi}{4} - \alpha = \phi' - \theta \Rightarrow \phi' = \frac{\pi}{4} - \alpha + \theta = \frac{\pi}{4} - \phi + 2\theta \tag{31}$$

It suffices to consider $k = 0$ if ϕ' is wrapped around to lie in the interval $\left[0, \frac{\pi}{2}\right)$.

REFERENCES

[1] B. Murmann, "ADC Performance Survey 1997–2015," [Online.] Available: http://www.stanford.edu/∼murmann/adcsurvey.html

[2] J. Singh, O. Dabeer, and U. Madhow, "On the limits of communication with low-precision analog-to-digital conversion at the receiver," *IEEE Trans. Commun.*, vol. 57, no. 12, pp. 3629–3639, Dec. 2009.

[3] J. Singh and U. Madhow, "Phase-quantized block noncoherent communication," *IEEE Trans. Commun.*, vol. 61, no. 7, pp. 2828–2839, Jul. 2013.

[4] D. A. Sobel and R. W. Brodersen, "A 1 Gb/s mixed-signal baseband analog front-end for a 60 GHz wireless receiver," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1281–1289, Apr. 2009.

[5] A. Host-Madsen and P. Handel, "Effects of sampling and quantization on single-tone frequency estimation," *IEEE Trans. Signal Process.*, vol. 48, no. 3, pp. 650–662, Mar. 2000.

[6] F. Sun, D. Liu, and G. Yue, "Particle filtering based automatic gain control for adc-limited communication," in *Proc. IEEE 73rd Veh. Technol. Conf. (VTC Spring)*, 2011, pp. 1–5.

[7] O. Dabeer and U. Madhow, "Channel estimation with low-precision analog-to-digital conversion," in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2010, pp. 1–6.

[8] A. Wadhwa, U. Madhow, and N. Shanbhag, "Space-time slicer architectures for analog-to-information conversion in channel equalizers," in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2014.

[9] O. Dabeer and E. Masry, "Multivariate signal parameter estimation under dependent noise from 1-bit dithered quantized data," *IEEE Trans. Inf. Theory*, vol. 54, no. 4, pp. 1637–1654, Apr. 2008.

[10] A. Wadhwa and U. Madhow, "Blind phase/frequency synchronization with low-precision adc: A bayesian approach," presented at the 51st Annu. Allerton Conf. Commun., Contr., Comput. (Allerton), 2013.

[11] D. Divsalar and M. K. Simon, "Multiple-symbol differential detection of mpsk," *IEEE Trans. Commun.*, vol. 38, no. 3, pp. 300–308, Mar. 1990.

[12] M. Naghshvar *et al.*, "Active sequential hypothesis testing," *Ann. Statist.*, vol. 41, no. 6, pp. 2703–2738, 2013.

[13] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Trans. Autom. Control*, vol. 58, no. 10, pp. 2451–2464, Oct. 2013.

[14] A. G. Busetto, A. Hauser, G. Krummenacher, M. Sunnåker, S. Dimopoulos, C. S. Ong, J. Stelling, and J. M. Buhmann, "Near-optimal experimental design for model selection in systems biology," *Bioinformatics*, vol. 29, no. 20, pp. 2625–2632, 2013.

[15] G. Atia and S. Aeron, "Asymptotic optimality results for controlled sequential estimation," presented at the 51st Annu. Allerton Conf. Commun., Contr., Comput. (Allerton).

[16] A. Krause and C. E. Guestrin, "Near-Optimal Nonmyopic Value of Information in Graphical Models," 2012, arXiv Preprint arXiv:1207.1394.

[17] F. Quitin, M. M. U. Rahman, R. Mudumbai, and U. Madhow, "A scalable architecture for distributed transmit beamforming with commodity radios: Design and proof of concept," *IEEE Trans. Wireless Commun.*, vol. 12, no. 3, pp. 1418–1428, Mar. 2013.

[18] "Characterization of Frequency and Phase Noise," Tech. Rep. 580 of the Int. Radio Consultative Committee (C.C.I.R.), pp. 142–150, 1986.

[19] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation: Theory Algorithms and Software*. New York, NY, USA: Wiley-Interscience, 2001.

[20] "HMC6001 Datasheet, Millimeter Wave Receiver," [Online.] Available: http://www.hittite.com/products/view.html/view/HMC6001

[21] D. Rife and R. R. Boorstyn, "Single tone parameter estimation from discrete-time observations," *IEEE Trans. Inf. Theory*, vol. IT-20, no. 5, pp. 591–598, Sep. 1974.

[22] S. Kullback, *Information Theory and Statistics*. Mineola, NY, USA: Courier Dover Publications, 1997.

Aseem Wadhwa received his bachelors degree in electrical engineering from the Indian Institute of Technology Delhi, in 2009, and masters and Ph.D. degrees in electrical and computer engineering from the University of California, Santa Barbara, in 2011 and 2014, respectively. He is interested in machine learning, neuroscience inspired learning and ADC constrained high speed communication. Aseem was a postdoctoral researcher with Prof. Upamanyu Madhow at UC Santa Barbara from 2015 to 2016, and currently works for Apple Inc, Cupertino.

Upamanyu Madhow is Professor of Electrical and Computer Engineering at the University of California, Santa Barbara. His research interests broadly span communications, signal processing and networking, with current emphasis on millimeter wave communication, and on distributed and bio-inspired approaches to networking and inference. He received his bachelor's degree in electrical engineering from the Indian Institute of Technology, Kanpur, in 1985, and his Ph.D. degree in electrical engineering from the University of Illinois, Urbana-Champaign in 1990. He has worked as a research scientist at Bell Communications Research, Morristown, NJ, and as a faculty at the University of Illinois, Urbana-Champaign. Dr. Madhow is a recipient of the 1996 NSF CAREER award, and co-recipient of the 2012 IEEE Marconi prize paper award in wireless communications. He has served as Associate Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS, the IEEE TRANSACTIONS ON INFORMATION THEORY, and the IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY. He is the author of two textbooks published by Cambridge University Press, *Fundamentals of Digital Communication* (2008) and *Introduction to Communication Systems* (2014).