

Blind phase/frequency synchronization with low-precision ADC: a Bayesian approach

Aseem Wadhwa and Upamanyu Madhow

Department of ECE, University of California Santa Barbara, CA 93106

Email: {aseem, madhow}@ece.ucsb.edu

Abstract—Modern communication receivers heavily leverage Moore’s law, which enables low-cost implementations of sophisticated functionalities in digital signal processing (DSP). However, as communication systems scale up in bandwidth, the availability of analog-to-digital converters (ADCs) becomes a fundamental bottleneck for such DSP-centric design. In this paper, we investigate a canonical problem of blind carrier phase and frequency synchronization in order to obtain insight into the performance limitations imposed by severe quantization constraints. We consider an ideal Nyquist sampled QPSK system with coarse phase quantization, implementable with one bit ADCs after analog linear combinations of in-phase (I) and quadrature (Q) components. We propose blind Bayesian algorithms for rapid phase acquisition, followed by continuous feedback-based phase/frequency tracking, based on jointly modeling the unknown phase and frequency, the unknown data, and the severe nonlinearity introduced due to coarse phase quantization. Our performance evaluation shows that excellent performance, close to that of an unquantized system, is achieved by the use of 12 phase bins (implementable using 6 one-bit ADCs).

I. INTRODUCTION

Modern communication transceiver designs leverage Moore’s law for low-cost implementation (e.g., for today’s WiFi and cellular systems), by using DSP to perform sophisticated functionalities such as synchronization, equalization, demodulation and decoding. The central assumption in such designs is that analog signals can be faithfully represented in the digital domain, typically using ADCs with 8-12 bits of precision. We would like to extend this approach to emerging communication systems employing bandwidths of multiple GHz, such as emerging millimeter wave wireless networks (e.g., using the 7 GHz of unlicensed spectrum in the 60 GHz band), as well as for signal processing in bandwidth efficient optical communication systems. The key bottleneck to doing this is the ADC: the cost and power consumption of high-precision ADCs become prohibitive at multi-GHz sampling rates [1]. Since we would like to continue taking advantage of Moore’s law despite this bottleneck, it is natural to ask whether DSP-centric architectures with samples

quantized at significantly less precision (e.g., 1-4 bits) can be effective. 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. In particular, we consider in this paper a canonical problem of blind carrier phase/frequency synchronization based on coarse *phase-only quantization* (implementable using digitally controlled linear analog preprocessing of I and Q samples, followed by one-bit ADCs), and develop and evaluate the performance of a Bayesian approach based on joint modeling of the unknown data, frequency and phase, and the known quantization non-linearity.

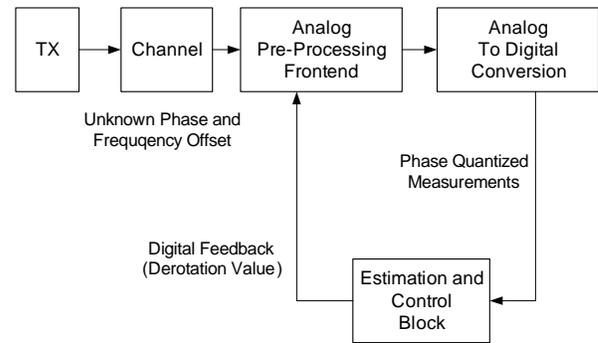


Fig. 1. Receiver Architecture

Receiver architecture: We consider differentially encoded QPSK over an AWGN channel. In order to develop fundamental insight into carrier synchronization, we do not model timing asynchronism or channel dispersion. 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 complex-valued symbol rate samples before passing it through the *ADC block*. The ADC block quantizes the phase of the samples into a small number of bins. Phase quantization (which suffices for hard decisions with PSK constellations) has the advantage of not requiring automatic gain control (AGC), since it can be implemented by passing linear combinations of the in-phase and quadrature components through one-bit ADCs (quantization into $2n$ phase bins requires n such linear combinations)

This research was supported in part by the Institute for Collaborative Biotechnologies through the 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.

[3]. 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. Design of this estimation and control block is the subject of this paper. We break the synchronization problem into two steps (a) rapid blind *acquisition* of initial frequency/phase estimates, (b) continuous *tracking* while performing data demodulation.

Contributions: For the acquisition step, we develop a Bayesian algorithm for blind phase estimation, which includes design of the feedback to the analog preprocessor to aid in estimation. The feedback evolves with the posterior distribution of the phase, and we show that an information-theoretically motivated greedy strategy is useful in improving performance at high SNR. Since frequency offsets between transmitter and receiver are typically much smaller than the symbol rate, the phase is well approximated as constant over multiple symbols, hence acquisition is performed ignoring frequency offset. For the tracking step, 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. 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.

Related work: A phase-quantized carrier-asynchronous system model similar to ours was studied in [4]. However, instead of explicit phase/frequency estimation and compensation as in this paper, block noncoherent demodulation, approximating the phase as constant over a block of symbols, was employed in [4]. Whereas a performance degradation of about 2 dB compared to the unquantized block noncoherent case was reported in [4], the algorithm proposed in this paper performs better, with bit error rates almost identical to the unquantized coherent system. Moreover, the analog preprocessing used in the tracking step is simpler compared to the dither scheme proposed in [4]. 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 [5], including blocks for both carrier synchronization and equalization. While the emphasis in [5] 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. Other related work on estimation using low-precision samples includes frequency estimation [6], amplitude estimation for PAM signaling [7],

channel estimation [8] and analysis of effects of quantization for fading channels [9][10].

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 k th sample. In order to simplify digital control of the derotation, we restrict the allowable *derotation values* θ to a finite set of values, denoted by \mathbb{C} ; in our simulations, we consider a phase resolution of the order of $2\pi/128$, which produces negligible degradation in coherent demodulation performance. After derotation, the sample is quantized into one of $M = 2n$ phase bins: $[(m-1)\frac{2\pi}{M}, m\frac{2\pi}{M}]$ for $m = 1, \dots, M$. In our simulations, we consider $M = 8$ and $M = 12$ (Figs. 3(a) and 4(a)). 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 ML decision boundaries for coherent QPSK demodulation.

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

$$z_k = Q_M \left(\arg \left(b_k e^{j(\phi_c + k \cdot 2\pi T_s \Delta f)} e^{-j\theta_k} + w_k \right) \right) \quad (1)$$

where,

- M := number of bins over $[0, 2\pi)$ for phase quantization;
- $z_k \in \{1, 2, \dots, M\}$ are the observations,
- $Q_M : [0, 2\pi) \rightarrow \{1, 2, \dots, M\}$ denotes the quantization function, $Q_M(x) = \lceil x \cdot \frac{M}{2\pi} \rceil$ for $x \in [0, 2\pi)$,
- $b_k \in \{e^{j\pi/4}, e^{j3\pi/4}, e^{j5\pi/4}, e^{j7\pi/4}\}$ normalized QPSK symbol transmitted, assumed to be uniformly distributed,
- $\phi_c, \Delta f$:= the unknown phase and frequency offset,
- T_s := symbol time period,
- $\theta_k \in \mathbb{C} = \{\text{mod}(i \cdot d\theta, 2\pi)\}, i \in \mathbb{I}$, the derotation value for the k th symbol, $d\theta$ denoting the phase resolution,
- w_k := independent complex AWGN, $\text{Re}(w_k) = \text{Im}(w_k) \sim \mathcal{N}(0, \sigma^2)$, where $\text{SNR per bit} = \frac{E_b}{N_0} = \frac{1}{2\sigma^2}$.

The carrier frequency offset Δf is typically of the order of 10-100 ppm of the carrier frequency. For example, for a 60 GHz link, the offset could be as large as 6 MHz, but is still orders of magnitude smaller than the symbol rate, which is of the order of Gsymbols/sec. Thus, it can be set to zero without loss of generality in the acquisition step (described in Section III), where we derive estimates of the unknown phase ϕ_c based on a small block of symbols. We do model the frequency offset in the tracking step (Section IV).

III. PHASE ACQUISITION

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

$$z_k = Q_M(u_k)$$

$$u_k = \arg \left(e^{jp_k \frac{\pi}{4}} e^{j\beta_k} + w_k \right) \quad (2)$$

$$\beta_k = \phi_c - \theta_k$$

where u_k denotes the unquantized phase, β_k captures the net rotation of the transmitted QPSK symbol and p_k 's are independent and uniformly distributed over $\{1, 3, 5, 7\}$, since we are interested in blind estimation (without the use of training symbols). We now drop the subscript k to simplify notation. Conditioned on β we can express the density of u as follows (derivation is presented in the appendix):

$$f_u(\alpha; \beta) = \sum_{i=1}^4 \frac{1}{4} f_{u|p=2i-1}(\alpha; \beta) \quad ; \quad \alpha \in [0, 2\pi)$$

$$f_u(\alpha; \beta) = \sum_{i=1}^4 \frac{1}{4} \left[\frac{a_i (2 - \operatorname{erfc}(\frac{a_i}{\sigma\sqrt{2}})) e^{\frac{a_i^2 - 1}{2\sigma^2}}}{2\sigma\sqrt{2\pi}} + \frac{e^{-\frac{1}{2\sigma^2}}}{2\pi} \right] \quad (3)$$

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

Looking at the expression above, if we define the density for $\beta = 0$ as $f_u(\alpha) = f_u(\alpha; 0)$, then the density at non-zero values of β can be evaluated simply by circular shifts (by 2π) of $f_u(\alpha)$. Another property to note is the periodicity of $f_u(\alpha)$ with period 90° (as shown in Fig. 2), which is due to the uniform distribution over the QPSK constellation. Conditional distribution of the quantized measurements can now be computed by evaluating the appropriate integrals:

$$P(z_k = m | \beta_k) = \int_{(m-1)\frac{2\pi}{M}}^{m\frac{2\pi}{M}} f_u(\alpha - \beta_k) d\alpha \quad (4)$$

$$\text{where } m \in \{1, 2, \dots, M\}$$

Using the expression above, given the k^{th} phase measurement in bin m , the single step likelihood of the phase is given by $l_k(\phi|m) = \log(p(z_k = m|\phi))$ if the derotation phase $\theta_k = 0^\circ$. Nonzero θ_k simply results in a circular shift of $l_k(\phi|m)$. Due to the periodicity of $f_u(\alpha)$, it suffices to limit ϕ to the interval $[0, 90^\circ]$. We drop the subscript k as noise is independent over symbols and $l(\phi|m) = l_k(\phi|m) \forall k$. The Bayesian estimator, as discussed next, essentially involves successively adding these single step likelihoods as more measurements are made. An interesting property to note is the periodicity of $l(\phi|m)$ in m with period $M/4$, which follows from the symmetry induced by equiprobability of the transmitted symbols. For example, if $M = 8$ (Fig. 3(a)), a measurement z_k in bin 1 or bin 3 results in the same likelihood function. Fig. 2 shows the three distinct likelihoods for $M = 12$ (6 one-bit ADCs).

A. Bayesian Estimation given Derotation Phases θ_k

Conditioned on the past derotation values θ_1^k (which are known) and the quantized phase observations z_1^k , applying

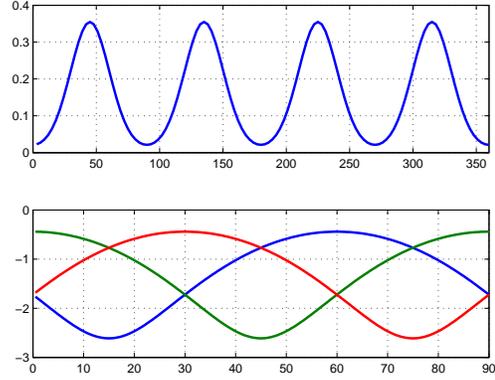


Fig. 2. **(top)** Probability Density of unquantized phase u at $\beta = 0$, $f_u(\alpha)$ **(bottom)** Single step likelihoods $l(\phi|m)$ given $z = m$ and $\theta = 0^\circ$ ($M = 12$, SNR=5dB). **blue:** $l(\phi|1) = l(\phi|4) = l(\phi|7) = l(\phi|10)$, **green:** $l(\phi|2) = l(\phi|5) = l(\phi|8)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

Bayes rule gives us a recursive equation for updating the posterior of the unknown phase as:

$$p(\phi|z_1^k, \theta_1^k) = \frac{p(z_k|\phi, \theta_k) p(\phi|z_1^{k-1}, \theta_1^{k-1})}{p(z_k|\theta_k)} \quad (5)$$

Normalizing the pdf obviates the need to evaluate the denominator. We now go to the log domain to obtain an additive update for the cumulative log likelihood. Denoting by $l_{1:k}(\phi) = \log(p(\phi|z_1^k, \theta_1^k))$ the cumulative update up to the k^{th} symbol, we update it recursively simply by adding the single step update $l_k(\phi) = \log(p(z_k|\phi, \theta_k))$, as follows:

$$l_{1:k}(\phi) = l_{1:k-1}(\phi) + l_k(\phi) \quad (6)$$

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

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

We start with a uniform prior $p(\phi)$ over $[0^\circ, 90^\circ)$. Single step likelihoods, $l(\phi|m)$ for $m = 1, \dots, M/4$, can be precomputed and stored offline, and circularly shifted by the derotation phase θ_k as the estimation proceeds. The recursive update (6) requires only the latest posterior to be stored.

B. Choosing the Derotation Phases θ_k

Setting the values of the derotation phases provides a means of applying a *controlled dither* prior to quantization. In this subsection, we investigate whether it could be used for *speeding up* the phase acquisition. We start by looking at two motivating scenarios where the naive strategy of setting $\theta_k = \text{constant} \forall k$ fails to give satisfactory results.

Example 1: Consider 8 phase quantization bins and $\phi_c = 10^\circ$ (Fig. 3). Choosing $\theta_k = 0^\circ \forall k$ results in a bimodal posterior with a spurious peak at $\phi = 35^\circ$. Due to symmetry of the phase boundaries and equiprobable distribution over the transmitted symbols, the set of observations (1,3,5,7) and

(2,4,6,8) leads to the posterior being updated in identical ways. With probability of getting bin 3 for $\phi = 35^\circ$ being equal to the probability of getting bin 1 for $\phi = 10^\circ$, there is an unresolvable ambiguity between the two phases. In general for any phase α , we have $P(z_k = i|\phi = \alpha, \theta_k = 0) = P(z_k = j|\phi = 45^\circ - \alpha, \theta_k = 0) \forall i, j \in \{1, 3, 5, 7\}$ or $\forall i, j \in \{2, 4, 6, 8\}$; which gives rise to a bimodal posterior with peaks at α and $45^\circ - \alpha$. Such ambiguities were also noted in the block noncoherent system considered in [3]. One approach to alleviate this ambiguity is to dither θ_k randomly; this dithers the spurious peak while preserving the true peak, leading to a unimodal distribution for the posterior 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 by the next example.

Example 2: Now consider 12 phase bins and no noise (or very high SNR), again with true phase offset $\phi_c = 10^\circ$. Since there is no noise, all observations fall in bins 2,5,8,11, resulting in a flat phase posterior over the interval $[75^\circ, 90^\circ] \cup [0^\circ, 15^\circ]$ if there is no dither ($\theta_k \equiv 0^\circ$). This could lead to an error as high as 25° (Fig. 4). On the other hand, using randomly dithered θ_k s 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.

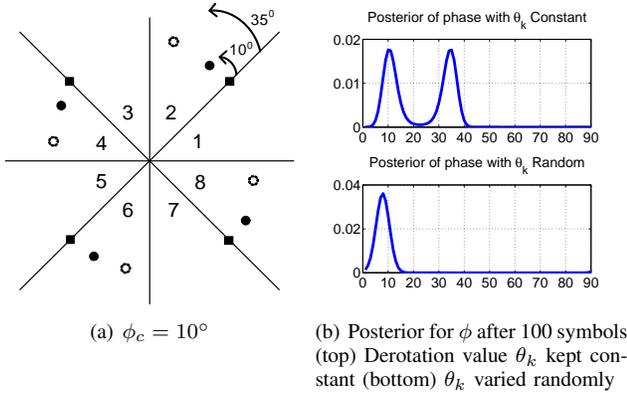


Fig. 3. Example 1: SNR=5dB, 8 uniform quantization regions

While randomly dithered derotation is a robust design which overcomes the shortcomings of the naive strategy of no dither, it is of interest to ask whether we can do better. Optimizing the sequence of derotation phases in order to minimize a performance criterion such as the mean squared error in the estimated phase is a difficult problem: even a genie-aided system which knows the true phase ϕ_c (which of course would obviate the need for phase estimation in the first place) leads to a computationally intractable Partially Observable Markov Decision Problem (POMDP). Instead, we propose an information-theoretically motivated *greedy entropy* strategy that minimizes the entropy of the posterior distribution across choices of derotation phase over the next step.

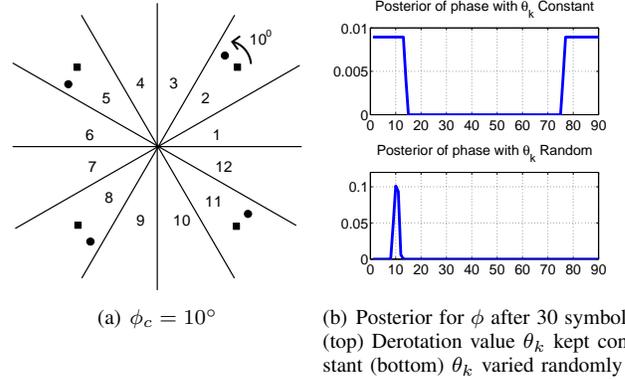


Fig. 4. Example 2: SNR=35dB, 12 uniform quantization regions

C. 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 $f_{k-1}(\phi) = p(\phi|z_1^{k-1}, \theta_1^{k-1})$. The uncertainty in this posterior, which is captured by its entropy, is a measure of the confidence in the MAP estimate ($\text{argmax}_\phi f_{k-1}(\phi)$) after $k - 1$ steps. For each possible action $\theta \in \mathbb{C}$ in the next step, we compute the *Expected Entropy* of the posterior after k^{th} step, denoted by $\bar{H}_k(\theta)$, as follows:

$$\bar{H}_k(\theta) = \sum_{m=1}^{M/4} P(z_k = m|\theta_k = \theta) \times H[p(\phi|z_1^{k-1}, \theta_1^{k-1}, \theta_k = \theta, z_k = m)] \quad (7)$$

where H denotes the entropy computed using a finely discretized version of the posterior,

$$p(z_k = m|\theta_k = \theta) = \sum_{\phi} p[z_k = m|\theta_k = \theta, \phi] f_{k-1}(\phi)$$

$$p(\phi|z_1^{k-1}, \theta_1^{k-1}, \theta_k = \theta, z_k = m) = c \cdot f_{k-1}(\phi) \cdot p(z_k = m|\phi, \theta)$$

Note that $\phi \in [0, \frac{\pi}{2}]$ and c is a normalizing constant (such that density sums to 1). The last equation follows from (5). Due to symmetry in probability of quantized phase measurements, we only need to sum over $M/4$ terms. The size of the set \mathbb{C} is limited by the resolution of the allowed derotation values $d\theta$ and size of the quantization bin $2\pi/M$ i.e. $|\mathbb{C}| = \lceil (2\pi/M)/d\theta \rceil$. The next action chosen is the one that leads to minimum expected future entropy:

$$\theta_k = \underset{\theta \in \mathbb{C}}{\text{argmin}} \bar{H}_k(\theta)$$

D. Simulation Results

The 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:

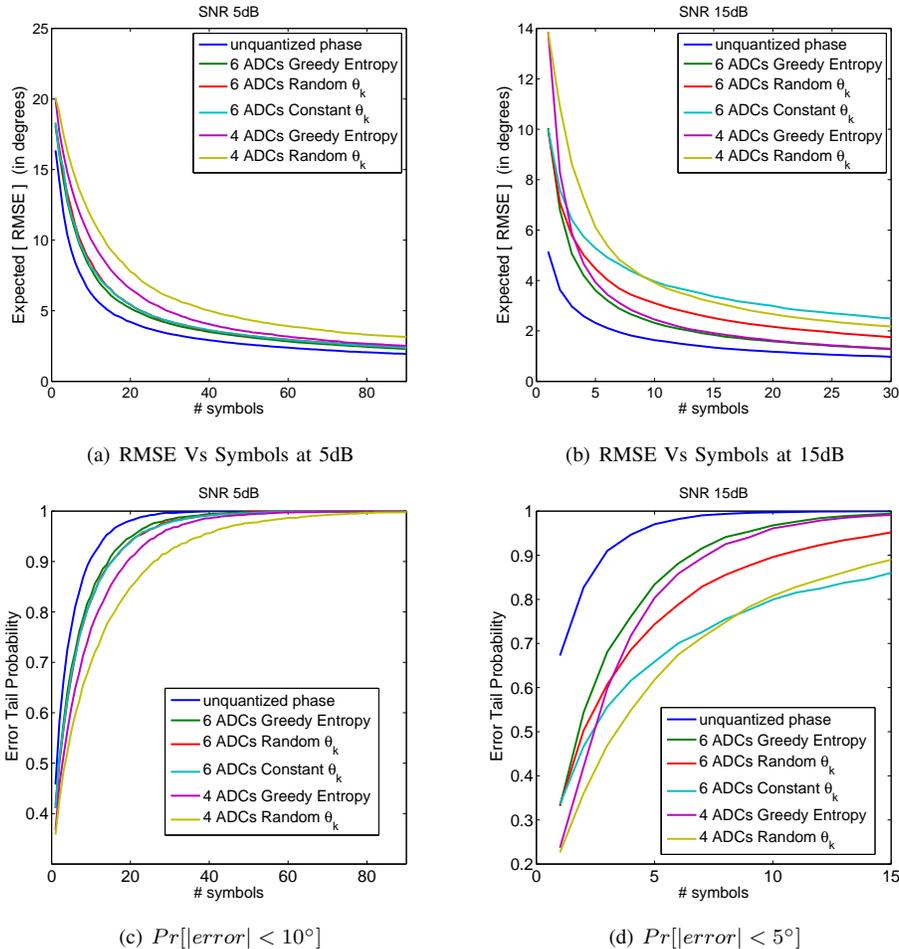


Fig. 5. Simulations of different strategies for choosing the feedback θ_k with 4 and 6 ADCs (8 and 12 phase bins)

a low value of 5 dB and a high value of 15 dB. The performance measures are the root mean squared error (RMSE), which captures average behavior, and the probability of the phase error being smaller than a threshold, which captures the tail behavior. We make the following observations: (a) Increasing the number of bins from 8 to 12 (4 ADCs to 6 ADCs) provides significant improvement, moving the curves significantly towards the unquantized limits. (b) The greedy entropy policy with 6 ADCs performs close to MAP estimation with unquantized observations, indicating it cannot be far away from the *optimal* control policy. (c) At high SNR, the naive policy of keeping the derotation values constant performs the worst, as expected. The greedy entropy policy drives the MSE to zero more quickly than random dithering. (d) At low SNR, there is little to distinguish between the different derotation policies for 6 ADCs, since the noise supplies enough dither to give a rich spread of measurements across different bins. However, when the quantization is more severe (4 ADCs), the greedy entropy policy provides performance gains over random dithering even at low SNR. To summarize, we find that efficient dithering policies could be very 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 aim to correct for the channel phase to enable accurate demodulation, in contrast to the acquisition phase, where the derotation is designed to aid in phase estimation.

IV. 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. The phase can be written as $\phi_c(k) = \phi_0 + 2\pi k T_s \Delta f = \phi_0 + k\eta$, where η is the *normalized frequency offset*, defined as the rate of change of phase in radians per symbol. To get a concrete idea of how fast the phase varies, consider the following typical values: $f_c = 60$ GHz, bandwidth of 6 GHz, i.e. $T_s = (6 \times 10^9)^{-1}$ secs, an offset $\Delta f = 100\text{ppm} \cdot f_c$, which leads to $\eta = 2\pi T_s \Delta f = 2\pi \cdot 10^{-3}$ radians; a linearly varying phase rate of 0.36° per symbol. We can therefore

accurately approximate the phase as roughly constant over a few tens of symbols, while obtaining an accurate estimate of the frequency offset η would require averaging over 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 previous section, with two key differences: the derotation phase value is our current best estimate of the phase, and we do not need to average over the possible symbols, since we operate in decision-directed mode. These relatively coarse phase estimates are then fed to an extended Kalman filter (EKF) for tracking both frequency and phase.

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

Process Model

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

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

where $w(k) \sim \mathcal{N}(0, Q_k)$ is the process noise, the state vector comprises the phase and the normalized frequency offset $x_k = [\phi(k) \ \eta(k)]^T$ and the state evolution matrix $A = [1 \ 1; 0 \ 1]$. Note that Q_k is of the form $\sigma_p^2 \cdot [1 \ 1; 1 \ 1]$ since the same noise term influences both the phase and frequency offset i.e. $\eta(k) = \eta(k-1) + w_k(2)$, and $\phi(k) = \phi(k-1) + \eta(k) = \phi(k-1) + \eta(k-1) + w_k(2)$, hence $w_k(1) = w_k(2)$.

Measurement Model

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

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

where $h(\cdot)$ is a non linear measurement function. The particular form is chosen to resolve the issue of unwrapping the phase periodically as it grows linearly: the factor of 4 inside the sine and cosine arguments chosen to obtain a period of 90° , since we are only interested in phase estimates over the range $[0, \pi/2]$. The measurement noise is $v(k) \sim \mathcal{N}(0, R_k)$. For the EKF, computation of the Jacobin 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 (we refer the readers to Chapter 10 of [11] for a discussion on EKF, and to [12] for a somewhat similar application of EKF for phase

tracking).

Time Update:

$$\hat{x}_{k|k-1} = A\hat{x}_{k-1}$$

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

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

Measurement Update:

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

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

\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 $\theta_k = \hat{x}_k(1) + \frac{W}{2} \cdot \hat{\eta}(k)$.

Tuning the filter: Although the measurement noise covariance R_k can be calculated from the variance of the posterior of the phase, constructed over the sliding window, the filter performance was observed to be quite robust to the choice of R_k over a range of SNR. For the simulations presented in this paper, we assumed a constant $R_k = [0.1 \ 0, 0 \ 0.1]^T$, which worked well for SNRs 0-15dB and sliding window length of $W = 50$ symbols. The scaling of the process noise (Q_k) trades off steady state versus tracking performance: small Q_k results in accurate estimates but slow reaction to abrupt changes in frequency, while large Q_k improves the response to abrupt changes at the expense of increased estimation error. Since the ultimate measure of performance is the bit error rate (BER) rather than the phase estimation error itself, a sensible approach to design is to set Q_k to the largest value (and hence the fastest response to abrupt changes) compatible with phase estimation errors causing a desired level of degradation in BER relative to ideal coherent demodulation.

A. Simulation Results

Fig. 6 shows the tracking algorithm in action. 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) η was changed from $2\pi \cdot 10^{-3}$ to $\pi \cdot 10^{-3}$ after 4000 symbols. The plot shows $\hat{\eta}$, the estimate, for choosing $Q_k = 5 \times 10^{-11}[1 \ 1; 1 \ 1]^T$ which enables the filter to lock onto the new value in about 1000 symbols. The last subplot 6(d) shows BER curves for ideal differentially decoded QPSK and that of the proposed algorithm, which is almost indistinguishable from the former. The slight loss of performance at high SNR is due to the assumption of finite resolution ($d\theta = \pi/64$) of the analog phase shifter in the front-end.

V. CONCLUSIONS AND FUTURE WORK

The framework for ADC-constrained receiver design illustrated in this paper has two core components:

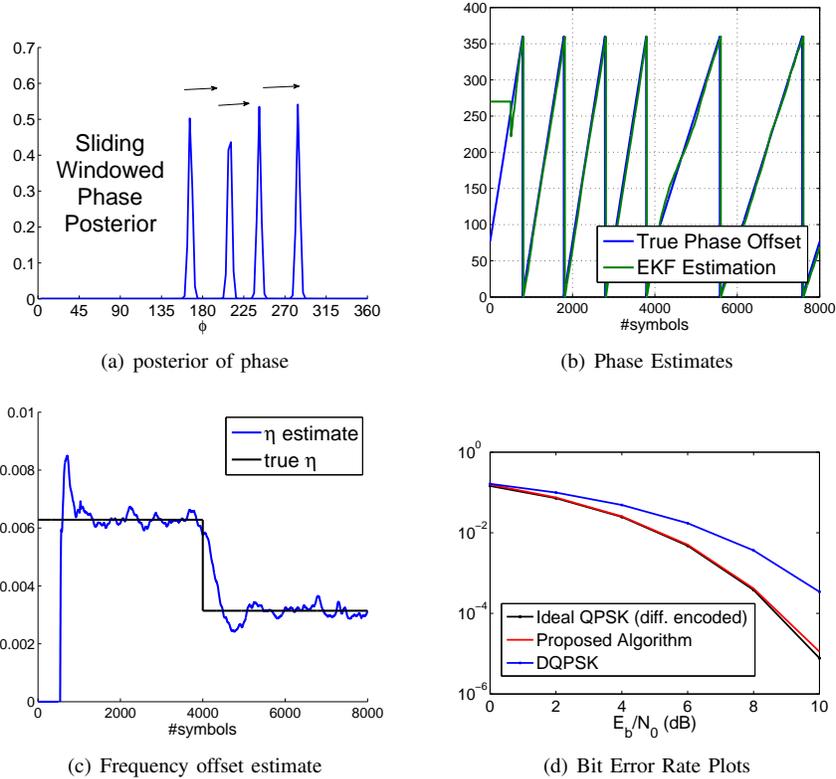


Fig. 6. Performance plots of EKF based Tracking Algorithm

(a) digitally controlled analog preprocessing: this provides the *dither* required for estimation with coarsely quantized observations in the acquisition step, and the *correction* required for coherent demodulation in the tracking step;

(b) Bayesian algorithms for estimation and feedback generation: this involves propagation of posterior probabilities in a manner that accounts for the quantization nonlinearity while probabilistically modeling unknown data and channel parameters. These posteriors are used to compute both the feedback for the analog preprocessor and the ultimate estimates of interest.

Our numerical results indicate that such architectures provide a promising approach for DSP-centric designs that exploit Moore's law despite the ADC bottleneck encountered at high communication bandwidths.

The success of a Bayesian approach for the simplified model considered here motivates future research on a comprehensive framework for receiver design subject to severe quantization constraints, addressing timing synchronization and dispersion compensation as well as carrier synchronization, and extending to larger amplitude/phase constellations. It is also of interest to develop a deeper theoretical understanding of fundamental performance limits under quantization constraints.

APPENDIX DERIVATION OF THE PHASE DISTRIBUTION

The expression for the unquantized phase is given by Eq. (2) as follows

$$u = \arg(e^{jp\frac{\pi}{4}}e^{j\beta} + w) = \arg(v)$$

p is uniformly distributed over $\{1, 3, 5, 7\}$ and w is complex WGN with variance σ^2 per dimension. Let us denote coordinates of the random complex variable v by $X = \text{Re}(v)$ and $Y = \text{Im}(v)$. Conditioned on p , $X \sim \mathcal{N}(\cos(p\frac{\pi}{4} + \beta), \sigma^2)$ and $Y \sim \mathcal{N}(\sin(p\frac{\pi}{4} + \beta), \sigma^2)$. To evaluate the distribution of the argument of v , we transform from cartesian to polar coordinates ($x = r\cos(\alpha)$, $y = r\sin(\alpha)$) which gives the following joint distribution

$$f(r, \alpha) = r^2 f(x, y)$$

$$f(r, \alpha) = \frac{r^2}{2\pi\sigma^2} e^{-\frac{1}{2\sigma^2}(x - \cos(p\frac{\pi}{4} + \beta))^2} e^{-\frac{1}{2\sigma^2}(y - \sin(p\frac{\pi}{4} + \beta))^2} \quad (8)$$

$$f(r, \alpha) = \frac{r}{2\pi\sigma^2} e^{-\frac{1}{2\sigma^2}(r^2 + 1 - 2r\cos(p\frac{\pi}{4} + \beta - \alpha))}$$

where (8) follows from the independence of X and Y . We can now marginalize out r to get the distribution of u

$$f_u(a) = \int_0^\infty \frac{r}{2\pi\sigma^2} e^{-\frac{1}{2\sigma^2}(r^2 + 1 - 2ra)} dr \quad (9)$$

$$a = \cos\left(p\frac{\pi}{4} + \beta - \alpha\right)$$

where dependence on α has been expressed through a . Integral (9) can be computed by observing that $f(a)$ (dropping subscript u) is the derivative of another integral $g(a)$ defined below, which in turn can be easily evaluated by completing squares in the exponent and expressing in terms of the standard Q function.

$$\begin{aligned} g(a) &= \frac{1}{2\pi} \int_0^\infty e^{-\frac{1}{2\sigma^2}(r^2+1-2ra)} dr \\ &= \frac{\sigma}{\sqrt{2\pi}} e^{-\frac{(1-a)^2}{2\sigma^2}} (1 - Q(a/\sigma)) \end{aligned}$$

$$f(a) = g'(a) = \frac{a(1 - Q(a/\sigma))e^{\frac{a^2-1}{2\sigma^2}}}{\sigma\sqrt{2\pi}} + \frac{e^{-\frac{1}{2\sigma^2}}}{2\pi} \quad (10)$$

Averaging out p we get Eq. (3).

REFERENCES

- [1] R. H. Walden, "Analog-to-digital converter survey and analysis," *Selected Areas in Communications, IEEE Journal on*, vol. 17, no. 4, pp. 539–550, 1999.
- [2] J. Singh, O. Dabeer, and U. Madhow, "On the limits of communication with low-precision analog-to-digital conversion at the receiver," *Communications, IEEE Transactions on*, vol. 57, no. 12, pp. 3629–3639, 2009.
- [3] J. Singh and U. Madhow, "On block noncoherent communication with low-precision phase quantization at the receiver," in *Information Theory, 2009. ISIT 2009. IEEE International Symposium on*, pp. 2199–2203, IEEE, 2009.
- [4] J. Singh and U. Madhow, "Phase-quantized block noncoherent communication," *CoRR*, vol. abs/1112.4811, 2011.
- [5] D. A. Sobel and R. W. Brodersen, "A 1 Gb/s mixed-signal baseband analog front-end for a 60 GHz wireless receiver," *Solid-State Circuits, IEEE Journal of*, vol. 44, no. 4, pp. 1281–1289, 2009.
- [6] A. Host-Madsen and P. Handel, "Effects of sampling and quantization on single-tone frequency estimation," *Signal Processing, IEEE Transactions on*, vol. 48, no. 3, pp. 650–662, 2000.
- [7] F. Sun, J. Singh, and U. Madhow, "Automatic gain control for ADC-limited communication," in *Global Telecommunications Conference (GLOBECOM 2010), 2010 IEEE*, pp. 1–5, IEEE, 2010.
- [8] O. Dabeer and U. Madhow, "Channel estimation with low-precision analog-to-digital conversion," in *Communications (ICC), 2010 IEEE International Conference on*, pp. 1–6, IEEE, 2010.
- [9] S. Krone and G. Fettweis, "Fading channels with 1-bit output quantization: Optimal modulation, ergodic capacity and outage probability," in *Information Theory Workshop (ITW), 2010 IEEE*, pp. 1–5, IEEE, 2010.
- [10] G. Middleton and A. Sabharwal, "On the impact of finite receiver resolution in fading channels," in *Allerton Conf. on Communication, Control and Computing*, 2006.
- [11] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation: theory algorithms and software*. Wiley-Interscience, 2001.
- [12] F. Quitin, M. M. U. Rahman, R. Mudumbai, and U. Madhow, "Distributed beamforming with software-defined radios: frequency synchronization and digital feedback," in *Global Telecommunications Conference (GLOBECOM 2012), 2012 IEEE*, (Anaheim, CA), IEEE, December 2012.