Frequency tracking with intermittent wrapped phase measurement using the Rao-Blackwellized particle filter

Maryam Eslami Rasekh *, Upamanyu Madhow*, Raghuraman Mudumbai †
*Department of Electrical and Computer Engineering, University of California Santa Barbara
Email: {merasekh, madhow}@ece.ucsb.edu
† Department of Electrical and Computer Engineering, University of Iowa
Email: rmudumbai@engineering.uiowa.edu

***Abstract*—In this paper, we consider the problem of frequency and phase tracking with intermittent measurements. While this problem is of fundamental interest, we are primarily motivated by the coherence requirements of distributed MIMO (DMIMO) applications in which cooperating nodes emulate a virtual antenna array and frequency synchronization among transmitting nodes is achieved by each node detecting and compensating for its frequency offset from a reference pilot that is broadcast periodically by a master node. Since accurate one-shot frequency estimation requires long measurement epochs, in the interest of minimizing overhead, we consider the problem of frequency/phase tracking using phase-only measurements. While the evolution of frequency and unwrapped phase can be modeled using a standard linear state space model, the measurement model is nonlinear, since the phase measurements are wrapped to a 2π interval. We develop a systematic design framework based on the concept of a Rao-Blackwellized particle filter: a particle filter is utilized to deduce the unwrapped phase from the wrapped phase measurements, while a Kalman filter estimates the frequency offset and provides phase predictions that are used to update particle weights. We show that the proposed filter effectively tracks the frequency and phase of on oscillator as long as the resync interval is smaller than a minimum that depends on the oscillator's drift statistics.**

I. Introduction

We consider the problem of frequency and phase tracking with intermittent measurements corrupted by additive white Gaussian noise. We are motivated by the tight synchronization requirements of distributed MIMO (DMIMO) applications in which cooperating nodes emulate a virtual antenna array (e.g., to form a beam or null towards a given destination). One approach to frequency synchronization in such a system is for a master node [1], or the destination node itself [2], to send a reference carrier that all cooperating nodes can lock on to. Given the drift inherent to oscillators, synchronization requires continuous tracking rather than just one-shot estimation. Our goal in this paper is to explore reduction of the overhead in such a procedure.

We consider the following canonical model, which focuses on a single node (e.g., any one of the cooperating nodes in a DMIMO system). The node receives bursts of a carrier waveform, corrupted by additive white Gaussian noise, over measurement epochs. Assuming that the frequency and phase of a carrier obeys a standard state space model, we ask how to reduce the measurement overhead, which can be accomplished by reducing the length of the measurement epochs and increasing the time between measurement epochs.

Well-known Cramer-Rao lower bounds [3], [4] on frequency and phase estimation over a single measurement epoch of length tell us that the variance of the phase error is inversely proportional to the SNR accumulated over the epoch, but that the rms frequency error is inversely proportional to the length T_m of the measurement epoch:

$$\sigma_\phi^2 \geq \frac{2}{SNR}$$

$$\sigma_f^2 \geq \frac{3}{2\pi^2 T_m^2 SNR}$$

Thus, as we shrink the length of the measurement epoch, the frequency estimate becomes unreliable. We therefore explore in this paper what can be accomplished with phase measurements alone (we assume, of course, that the integrated SNR over the measurement epoch is large enough for sufficiently accurate phase measurements).

If the unwrapped phase were available, then the standard Kalman filter provides a straightforward approach for tracking frequency and phase using a standard state space model. The key hurdle here is that we only have access to the wrapped phase, but wish to track frequency and phase despite large initial uncertainties in frequency, in a manner robust to jumps in frequency and phase. A fundamental difficulty is the *frequency aliasing* due to intermittent measurements of the wrapped phase: based on measurements spaced by T_s, we cannot distinguish between frequencies spaced by integer multiples of $1/T_s$.

There are three key components to our approach. The first is the use of a particle filter to explore multiple hypotheses regarding the unwrapped phase. The second component is the use of *Rao-Blackwellization,* or using the state space structure of the system for each particle to improve estimation performance, thus greatly reducing the number of particles required. That is, for each particle, it becomes possible to use the Kalman filter based on the unwrapped phase hypothesis of that particle to track frequency and phase. The third key component is dithering of the intervals between measurement epochs to resolve the problem of aliasing. We provide a deterministic dithering scheme that successfully resolves this ambiguity far more quickly than randomized dithering.

While particle filters have been employed before for pilot-based frequency offset estimation [5], [6], to the best of our knowledge, this is the first paper to show that it is possible to successfully track both frequency and phase based on wrapped phase measurements.

II. Model and Algorithm

The state space model used for the phase and frequency offset of a transmitter relative to the receiver in the discrete feedback times is as follows:

$$x_t = \mathbf{F} x_{t-1} + \nu_t$$

where the state $x_t = [\phi_t, \omega_t]^T$ is the relative phase and angular frequency offset of the transmitter and the state update matrix, $\mathbf{F}$, is defined by

$$\mathbf{F} = \begin{bmatrix} 1 & T_s \\ 0 & 1 \end{bmatrix}$$

and T_s is the period of the feedback samples. The process noise, ν_t, models the relative phase and frequency drift of the oscillators and has the distribution defined in (1) taken from the model introduced in [7], [8].

$$\nu_t \sim N(0, Q)$$

$$Q = \omega_c^2 q_1^2 \begin{bmatrix} T_s & 0 \\ 0 & 0 \end{bmatrix} + \omega_c^2 q_2^2 \begin{bmatrix} \frac{T_s^3}{3} & \frac{T_s^2}{2} \\ \frac{T_s^2}{2} & T_s \end{bmatrix} \quad (1)$$

q_1^2 and q_2^2 are the phase and frequency drift parameters and for the USRP oscillators used in this analysis have values $q_1^2 = 8.47 \times 10^{-22}$ s and $q_2^2 = 5.51 \times 10^{-18}$ Hz as deduced from the Allan variance measurements according to [8]. The measurement model is described in (2) and provides us with a measurement of the *wrapped* phase. The wrapped phase measurement noise distribution is approximated as Gaussian and a standard Kalman filter is used to track the state.

$$\begin{bmatrix} i_t \\ q_t \end{bmatrix} = \mathbf{h}(x_t) = A \begin{bmatrix} \cos(\phi_t) \\ \sin(\phi_t) \end{bmatrix} + \mathbf{N}_t, \quad \mathbf{N}_t \sim N(0, \sigma^2 I)$$

$$y_t = \angle(i_t + jq_t) = \begin{bmatrix} 1 & 0 \end{bmatrix} x_t + e_t, \quad e_t \sim N(0, r) \quad (2)$$

In order to make accurate frequency offset estimations, the unwrapped phase needs to be estimated. However, it is clear from (2) that a Kalman filter can only produce wrapped phase measurements. Therefore a particle filter is used to pinpoint the number of full 2π rotations of the unwrapped phase and by adding this amount to the wrapped phase an estimate for unwrapped phase is obtained. This leads to a Rao-Blackwellized particle filter (RBPF) framework where the intractable part of the state is tracked using a particle filter, complemented by a Kalman filter that is used for the linear portion of the state. The advantage of combining the two methods is that the dimension of the state that needs to be approximated by the particle filter is decreased allowing for far fewer particles required to cover the entire state space. The reader is referred to [9] for a more detailed description of the RBPF.

Different oscillators with the same nominal frequency can have tens of ppm of relative frequency offset, therefore to resolve this ambiguity, a set of N particles spanning the ambiguous bandwidth is considered. Each particle has a hypothesis about the number of full 2π rotations in the time interval T_s. A Kalman filter attached to each particle tracks the linear state corresponding to that particle. Based on the observed error in the filter, particle weights are assigned and those with high error are discarded until only the correct hypothesis remains. Once the unwrapped phase is obtained, as long as the frequency/phase drift between consecutive feedback bursts is significantly smaller than a full rotation and SNR is high enough for useful measurements, the single remaining Kalman filter will track the state correctly with high probability. If, due to sudden jumps in oscillator frequencies or temporary loss of the feedback signal, convergence is lost, the KF will show high error in its output for many of the iterations and convergence loss can be detected. By re-initiating a sufficient number of particles around the last estimate and running the particle filter, synchronization can be restored. Usually, in this case, the particles need to cover a smaller frequency span than during initialization and fewer particles are required. The formulation of this implementation is given in the appendix.

III. Frequency Aliasing

In the noise free case, the frequency estimate from two wrapped phase measurements, ϕ_0 and ϕ_1, that are t_1 seconds apart can be any of the values $(2\pi k + \phi_1 - \phi_0)/t_1$ where $k \in \mathbb{Z}$ and each value corresponds to a possibility for the unwrapped phase change in the period of t_1 (Fig. (1)). Repeating the procedure with a different amount of time between the phase samples, t_2, produces a similar set of possibilities for the frequency with a different distribution and the correct frequency is among the common estimates of the two sets. By continually changing the time between phase measurements, the correct frequency can be singled out as the common point of all the resulting sets.

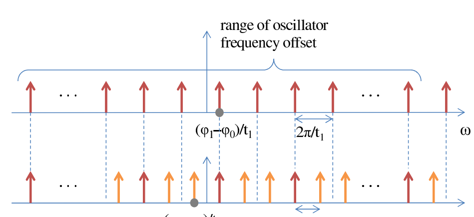

Fig. 1. Resolving frequency aliasing using non-uniform time sampling of wrapped phase

In the RBPF framework, this is done by assigning each frequency estimate of the first measurement interval to a particle and changing the time between phase measurements to eliminate all the incorrect estimates after a number of sampling cycles. Assuming the first feedback period is T_s, subsequent periods are increased (or decreased) by a factor of Δ from the initial value. Each particle is paired with a KF that updates and predicts the state at each measurement. Incorrect particles will produce high prediction errors in some of the time steps while the correct particle predicts the measurements with small error every time.

To see how much dithering of T_s is required, consider the ith particle. Without loss of generality we assume that particle $m = 0$ has the correct hypothesis and therefore the ith particle is $(i - m)2\pi/T_s = 2\pi i/T_s$ off in its frequency estimate. By increasing the time between two samples to $(1 + \Delta)T_s$, the phase prediction of the ith particle is

$$\frac{2\pi i + \phi_0}{T_s} \times (1 + \Delta)T_s$$

where ϕ_0 is the phase difference between the first two measurements. Then the difference between the (wrapped) phase predicted by this particle and the correct phase prediction is

$$\frac{2\pi i}{T_s} \times (1 + \Delta)T_s\Big]_{-\pi}^{\pi} = 2\pi i\Delta]_{-\pi}^{\pi} \tag{3}$$

Note that wrapped phase measurements allow for a maximum detectable error of π radians. Assuming the noise is not too high, a deterministic dithering strategy can be used to allow for fast convergence to the correct unwrapped phase. Random dithering techniques can also be used, but discussion of such schemes is omitted due to lack of space.

The deterministic dithering scheme we consider is as follows. We cycle through Δ from $\frac{1}{2}, \frac{1}{2^2}, ..., \frac{1}{2^L}$, where $L = \lceil 1 + \log_2 N \rceil$ is the number of steps needed to cover N different hypotheses for unwrapped phase (including the correct estimate) separated by steps of 2π.

$$\Delta(t) = (1 - \delta(\tilde{t}))\frac{1}{2^{\tilde{t}}} \tag{4}$$
$$\tilde{t} = \mathrm{mod}(t, L)$$

where $\delta(.)$ is the discrete-time unit impulse function. The following lemma states that, in the absence of noise, one such cycle is enough to reject all incorrect particles. (In a noisy setting, the weights for the incorrect particles are attenuated.)

Lemma 1. *The deterministic dither scheme (4) discards all incorrect particles within a cycle, while providing the highest possible noise margin of π for each particle in at least one step.*

Proof. From (3), for a frequency offset of $2\pi i/T_s$, the dither $\Delta = \frac{1}{2}$ sends the phase offset for all odd i to π. In general, $\Delta = \frac{1}{2^m}$ sends the phase offset for all odd $i/2^m$ to π. Thus, for $i = 1, ..., N \leq 2^L$ we take at most L iterations to cover all i in this fashion. □

IV. Implementation and Performance

The procedure is initialized by creating N particles each assuming a number of full 2π rotations in the interval of T_s ranging from $-(N - 1)/2$ to $(N - 1)/2$. The particles therefore cover a frequency span of $2\pi N/T_s$, with N chosen large enough so that this value covers the frequency span of oscillators' offset from their nominal frequency. After the first two measurements that are T_s seconds apart, the wrapped phase change estimate is added to the number of rotations to form the unwrapped phase change estimates and the frequency estimates are updated accordingly. In order to distinguish incorrect particles, the dithering scheme of (4) is used. Fig. 2 shows the progression of the weights of different particles in a sample simulation. At each step, half of the particles are severely attenuated and their weight is spread out on the remaining particles during normalization. The number of particles in this simulation was 126. It can be seen that all incorrect particles die down by step $1 + \log_2 126 \leq 8$.

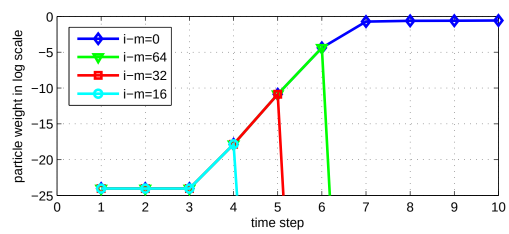

Fig. 2. Progression of particle weights with time steps

After convergence to the correct unwrapped phase, the frequency estimation error is that of a KF, as long as there are no cycle slips or loss of convergence. The criteria for maintaining convergence comes from the fact that the wrapped phase measurements can only correctly detect a phase change from one cycle to the next that is less than π. This means that the combination of process noise and measurement noise for the phase should have variance much smaller than π to ensure that noise samples will be smaller than π with near 1 probability. From (1) and (2) the total phase noise variance is observed to be:

$$\sigma_\phi^2 = \omega_c^2 q_1^2 T_s + \omega_c^2 q_2^2 \frac{T_s^3}{3} + r$$

where r is the phase measurement noise. Thus, the requirement for maintaining convergence with regard to the process and measurement characteristics and feedback period can be stated as

$$\sigma_\phi^2 = \omega_c^2 q_1^2 T_s + \omega_c^2 q_2^2 \frac{T_s^3}{3} + r \ll \pi^2 \tag{5}$$

Note that while our algorithm correctly estimates and tracks the unwrapped phase as long as the phase drift between successive feedback packets is (significantly) smaller than π, we may need more frequent feedback in order to maintain phase coherence sufficiently to achieve significant beamforming gains in distributed MIMO settings (e.g., we must keep phase errors

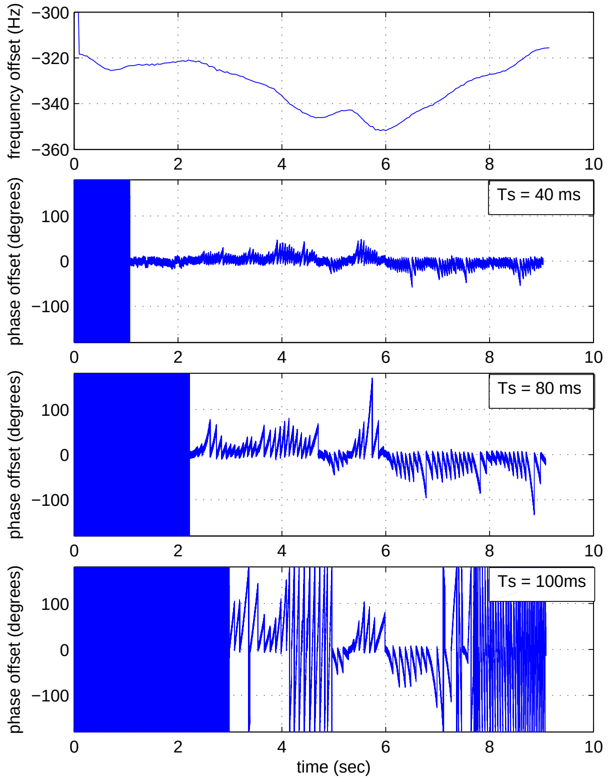

Fig. 3. Phase error of reconstructed signal for different values of feedback period (USRP oscillator)

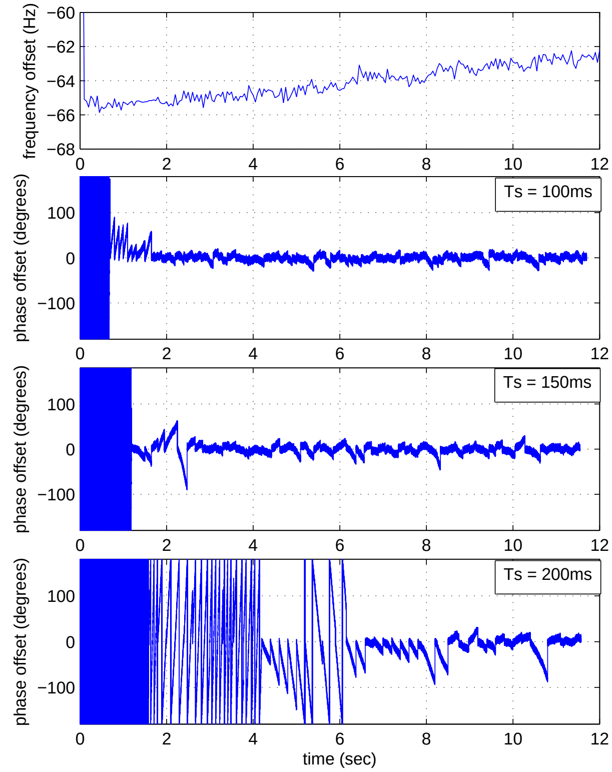

Fig. 4. Phase error of reconstructed signal for different values of feedback period (TCXO)

to within $\pi/12$ to attain 95% of the ideal beamforming gain [1]).

The proposed method was used to track the frequency of a USRP oscillator [10] over a 10 second interval of measured data. The program uses one phase measurement sample in each feedback cycle to estimate the frequency offset and reconstructs the signal using the estimations. The phase difference between the reconstructed signal and the actual signal is depicted in Fig. 3 along with the frequency drift of the oscillator. The deterministic dithering scheme was used and the results are shown for different values of T_s. As expected from (5), increasing T_s increases the chances of convergence loss since the drift causes phase errors that are too large to track using wrapped phase measurements. Performance degradation is observed around $T_s = 80$ms or $T_s(1+\Delta) = 120$ms which according to (5) corresponds to a (maximum) phase error variance of around 20° which is not small enough compared to 180° to prevent convergence loss.

The same results are shown in Fig. 4 for the TCXO oscillators [11] that have a more stable carrier frequency and better drift parameters as seen in the frequency vs. time plot of Fig. 4. As expected, these oscillators can be tracked with larger T_s before the synchronization breaks down due to drift and convergence is lost. In Fig. 5 an instance of convergence loss due to oscillator frequency jumps has been displayed and the detection and recovery from convergence loss is visible in the phase error of the reconstructed signal.

To evaluate the statement of (5), simulated data was used with known q parameters and negligible measurement noise. Fig. 6 shows the regions of convergence in the q_2^2 - T_s plane when q_1^2 is kept fixed at the level calculated for the USRP oscillators and q_2^2 is normalized to the value obtained for the same oscillators. The value of σ_ϕ^2 obtained from (5) has been plotted on the same plane and the correlation between convergence and σ_ϕ^2 is seen clearly. Similar results are seen in the q_1^2 - T_s plane and the q_2^2 - q_1^2 planes but the results have

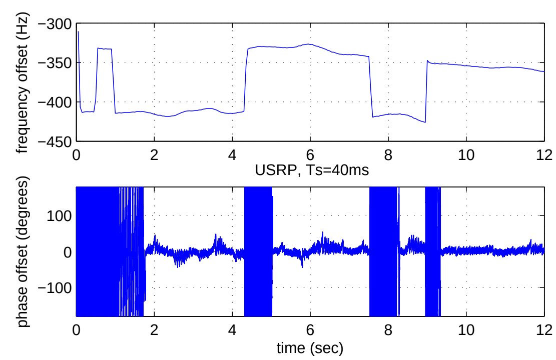

Fig. 5. Recovery from convergence loss due to sudden jumps in oscillator frequency

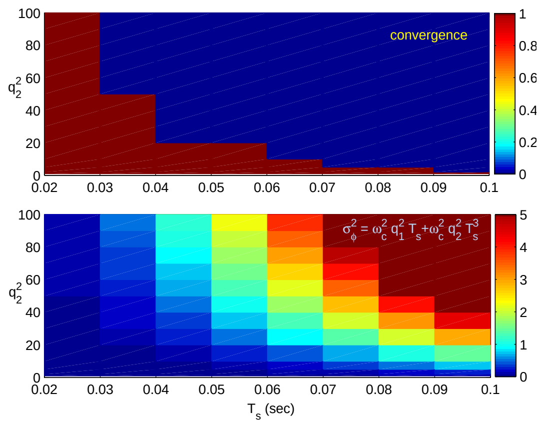

Fig. 6. Convergence and σ_ϕ^2 as a function of q_2^2 and T_s

been omitted due to space constraints.

V. Conclusions

We have presented a frequency tracking method using quasi-periodic measurements of wrapped phase for frequency synchronization in the distributed transmit beamforming framework. The proposed method combines a particle filter with a Kalman filter to obtain the unwrapped phase offset between the destination oscillator and the LO at the transmitter. The unwrapped phase is used to estimate the frequency offset of the distant oscillator with the high accuracy required for distributed transmit beamforming. The method performs well when the phase prediction error caused by measurement noise and oscillator drift in each cycle is low enough to be detected by the wrapped phase measurements, thus producing a limit on the feedback period. Results of the tracking method were presented for USRP oscillators with relatively high drift as well as the more stable TCXO oscillators. In both cases, for low enough values of feedback period the algorithm converged to the correct frequency in less than ten feedback cycles and phase error was maintained low enough to sustain the desired combination gain throughout the 10 second interval of measurement data.

Acknowledgments

This research was supported in part by the National Science Foundation under grants CCF-1302114, CCF-0830747, EPS-1101284, CAREER ECCS-1150801, CNS-1329657 and CCF-1302456, ONR grant N00014-13-1-0202 and a grant from the Army Research Lab's MAST program

Appendix

The system is set up with $N = 2L + 1$ particles that are initialized after two subsequent measurements that are T_s seconds apart as:

$$\phi^i_{0|0} = 2\pi i + [y_0 - y_{-1}]^{\pi}_{-\pi}, \quad \omega^i_{0|0} = \frac{\phi^i_{0|0}}{T_s}, \quad P_{0|0} = \mathbf{0}$$

where $i = -L, ..., L$. The prediction stage is formulated as:

$$\hat{x}_{t|t-1} = \mathbf{F}_t \hat{x}_{t-1|t-1}, \qquad P_{t|t-1} = \mathbf{F}_t P_{t-1|t-1} \mathbf{F}_t^T + Q_t$$

where $\mathbf{F}_t$ and Q_t are defined in section II and depend on t due to the dithering of T_s. Upon each wrapped phase measurement, y_t, particle weights are updated based on their prior probability for that observation, which follows a Gaussian distribution with zero mean and covariance P_ϕ:

$$W^i_t = W^i_{t-1} \frac{1}{\sqrt{2\pi P_\phi}} exp(-\frac{1}{2P_\phi} z_t^2)$$

$$z_t = y_t - H\hat{x}_{t|t-1} = (y_t - \hat{\phi}_{t|t-1})\Big]^{\pi}_{-\pi}, \quad P_\phi = P_{t|t-1}(1,1)$$

where $H = \begin{bmatrix} 1 & 0 \end{bmatrix}$. The Kalman estimate is updated as:

$$\hat{x}_{t|t} = \hat{x}_{t|t-1} + K_t z_t, \qquad P_{t|t} = (I - K_t H) P_{t|t-1}$$

$$K_t = P_{t|t-1} H^T (H P_{t|t-1} H^T + r)^{-1}$$

Particles with weights below a threshold are discarded and the weights are normalized across particles. The value of this threshold is chosen according to system noise levels to prevent discarding of the correct particle.

References

[1] R. Mudumbai, G. Barriac, and U. Madhow, "On the feasibility of distributed beamforming in wireless networks," *Wireless Communications, IEEE Transactions on*, vol. 6, no. 5, pp. 1754–1763, 2007.

[2] 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," *Wireless Communications, IEEE Transactions on*, vol. 12, no. 3, pp. 1418–1428, 2013.

[3] D. Rife and R. Boorstyn, "Single tone parameter estimation from discrete-time observations," *Information Theory, IEEE Transactions on*, vol. 20, no. 5, pp. 591–598, 1974.

[4] S. Kay, "A fast and accurate single frequency estimator," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 37, no. 12, pp. 1987–1990, 1989.

[5] K. J. Kim, M.-O. Pun, and R. Iltis, "Joint carrier frequency offset and channel estimation for uplink mimo-ofdma systems using parallel schmidt rao-blackwellized particle filters," *Communications, IEEE Transactions on*, vol. 58, no. 9, pp. 2697–2708, September 2010.

[6] W. Ng, C. Ji, W.-K. Ma, and H. So, "A study on particle filters for single-tone frequency tracking," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 45, no. 3, pp. 1111–1125, July 2009.

[7] D. Brown, P. Bidigare, and U. Madhow, "Receiver-coordinated distributed transmit beamforming with kinematic tracking," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, 2012, pp. 5209–5212.

[8] C. Zucca and P. Tavella, "The clock model and its relationship with the allan and related variances," *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 52, no. 2, pp. 289–296, 2005.

[9] F. Lindsten, "Particle filters and markov chains for learning of dynamical systems," Ph.D. dissertation, Linkping UniversityLinkping University, Automatic Control, The Institute of Technology, 2013.

[10] "Usrp products." [Online]. Available: http://ettus.com/

[11] Crystek, "Cpro33 (pocket reference oscillator)." [Online]. Available: http://www.crystek.com/home/pro/cpro33.aspx