Stochastic Modeling of Phase Noise in Distributed Coherent Passive Radar Systems

Andrew N. Morabito, John D. Sahr,
Zac M.P. Berkowitz, Laura E. Vertatschitsch
University of Washington Electrical Engineering,
Seattle WA 98195-2500, USA
morabito@u.washington.edu

Abstract

Passive radars sense targets by leveraging power from noncooperative transmitters. For coherent processing, highly stable oscillators must be used to drive each receiver of a distributed system. These oscillators, however, inject some amount of phase noise into the system as they are never perfectly stable. This phase noise degrades the Doppler resolvability of the radar. We will offer a technique to extract estimates of the phase noise process through naturally collected data from an example radar system. We will then fit stochastic models to the data and examine how well these models hold over time.

I. Introduction

Most literature dealing with modeling phase noise attempts to characterize the frequency-domain structure of the process for use in oscillator design [5], [3], [6]. Time-domain analysis techniques involve simulating the power spectra of the phase noise processes by modeling the time dynamics of the oscillator circuit [9]. In our technique we extract estimates of the phase noise time series from raw radar data and attempt to fit stochastic models to that data, characterizing the process without making assumptions about the structure of the oscillator.

Stochastic models for phase noise have been used in simulation before. Both autoregressive (AR) [1] and Wiener process [11], [10] models have been used, and we will attempt to fit these models to our data. We will denote the general class of autoregressive moving average processes (ARMA) to have AR parameters, ϕ_k's, and moving average (MA) parameters, θ_k's. We will consider a Wiener process to be parameterized by σ^2, where the process has mean zero and variance at time t equal to $t\sigma^2$.

By modeling the data we may obtain a better understanding of our system. If the model parameters can be related to system performance, then recurring estimation of those parameters may be new metrics worth reporting to indicated radar performance status. As our method will be treating the total phase noise accumulation as a single process to be modeled, the estimates will be representative of both the performance of the system oscillators as well as any other phase noise in the system, such as the possibility of phase noise due to atmospheric effects.

Our analysis of phase noise will be developed through its role in the Manastash Ridge Radar (MRR) [7]. MRR is a coherent, distributed passive radar system based at the University of Washington. A receiver at the University of Washington listens to FM broadcasts, recording them as a reference signal. A second receiver located at the Manastash Ridge Observatory (MRO), near Ellensburg WA, collects the scattered signals for use in correlation based processing. MRR has the added geographic advantage that the Cascade Mountain Range, which contain Mt. Rainier as the largest peak, suppress the direct path transmitter power. The reference and scatter signals are time identified with GPS receivers for use in processing. Each receiver contains a Phase Locked Loop (PLL) [1] that converts a high quality, 10 MHz GPS reference signal to the 56 MHz signal needed by the digitizer card. This PLL also contributes the main source of phase noise in the system.

II. Phase Noise Extraction

Using knowledge of MRR's system topology, we can extract estimates of the phase noise process and perform time series analysis. We mentioned that Mt. Rainier and the Cascade Mountain Range provide a natural barrier to most of the direct path energy emitted from Seattle FM broadcasts. In fact, Mt. Rainier scatters a substantial signal into the MRO receivers. Since Mt. Rainier is not moving, the scattered signal is not affected by a Doppler shift and is purely a delayed version of the transmitted signal (we will denote the range lag corresponding to this delay as r_0) at a lower power level and affected by a separate phase noise process. In this analysis, we will also consider

[1]MRR currently uses the NovaSource NS2-0045-025 PLL from Nova Engineering in its receivers.

TABLE I
COMPARISON OF MODEL FITS FOR AR PROCESSES FOR A GIVEN DATA SET

p	0	1	2	3	4	5	6	7	8	9	10
θ_1	-0.332	-0.939	-0.943	-0.938	-0.934	-0.925	-0.922	-0.934	-0.938	-0.930	-0.918
AIC	45067	43004	42972	42921	42910	42849	42846	42771	42758	42719	42652

that collection of the transmitted and scattered signals may each be affected by additive zero mean noise processes. These signals will be defined as:

$$x_\epsilon[n] = x_0[n]e^{j\phi_x[n]} + \epsilon_x[n] \tag{1}$$

$$y_\epsilon[n] = \alpha x_0[n-r_0]e^{j\phi_y[n]} + \epsilon_y[n] \tag{2}$$

We can then compute the detected signal at a range corresponding to Mt. Rainier:

$$\begin{aligned} a_\epsilon[n, r_0] &= \left(x_0^*[n]e^{-j\phi_x[n]} + \epsilon_x^*[n]\right)\left(\alpha x_0[n]e^{j\phi_y[n+r_0]} + \epsilon_y[n+r_0]\right) & (3) \\ &= \alpha\left|x_0[n]\right|^2 e^{j(\phi_y[n+r_0]-\phi_x[n])} + x_0^*[n]e^{-j\phi_x[n]}\epsilon_y[n+r_0] \\ &\quad +\alpha x_0[n]e^{j\phi_y[n+r_0]}\epsilon_x^*[n] + \epsilon_x^*[n]\epsilon_y[n+r_0] & (4) \end{aligned}$$

By then performing a coherent integration step the last three terms of Equation 4 are driven toward zero since x_0 and the additive noise processes are all mean zero and jointly uncorrelated. Examining the phase of $a_\epsilon[n, r_0]$, we obtain an estimate of $\phi_e \equiv \phi_y[n+r_0] - \phi_x[n]$ that we can use for time series analysis.

We can determine r_0 by finding the range lag that reflects the most power to the scatter signal receiver. In practice, finding the phase of $a_\epsilon[n, r_0]$ is challenging as we are interested in the unwrapped version of the principal value of the phase process. We recognized that instead of extracting ϕ_e (which requires phase unwrapping), we could perform analysis on another time series, ϕ_Δ, which is equal to the first difference on ϕ_e. This can be estimated simply by taking the first difference of the principal value of the phase from a_ϵ.

Fortunately, we can still analyze both of the proposed time series models using the differenced time series. A Wiener process model is simply analyzed since the time series is already differenced. We estimate the variance and use that to parameterize a Wiener process model. We can also detect AR(p) processes within this differenced time series since the first difference of an AR(p) process is an ARMA(p,1) process with AR coefficients equal to those of the AR(p) process and an MA coefficient of -1. Therefore if we model ϕ_Δ as ARMA(p,1), it corresponds to a model of ϕ_ϵ that is "AR-like". The closer θ_1 is to -1, the more "AR-like" ϕ_ϵ is.

III. STOCHASTIC MODELING

A. AR Process

We will begin by modeling the data as an AR process. In order to find an AR(p) model for the phase noise process, we must find an ARMA(p,1) model for the differenced time series. We will use the open source statistical software package R for this purpose [2]. For a given ARMA order, R produces the estimate of the parameters and the model's Akaike Information Criterion (AIC). As lower values of AIC correspond to better models, we are looking for models that produce the minimal value of AIC.

Table I shows a comparison of ARMA(p,1) models. We see a trend that the AIC appears to decrease as the order increases. We also see that the estimate of θ_1 does not appear to be exactly equal to -1 for any order we have tried to fit. Performing this technique over many data sets has shown that there has generally been a contiguous range of p values where the estimate of θ_1 has been reasonably close to -1 and deviates far from it otherwise. In this case, that range appears to begin at $p = 1$ and ends at some value $p > 10$. We are not searching that entire space to see the other end of the range since we have also found that the value closest to -1 tends to appear on the low side of the range. Since in this model selection we are concerned with fitting an ARMA(p,1) process that, when integrated, corresponds to an AR(p) process, we choose the model order that has a θ_1 closest to -1. In this case, that means selecting $p = 2$. Since the θ_1 estimate never quite reaches -1 (and the standard deviation of the estimate is generally on the order of .005), the phase process we are modeling is probably not exactly AR, but this appears to be our best approach in modeling it as an AR process.

[2]available at http://cran.r-project.org

TABLE II
TIME SERIES ANALYSIS RESULTS FOR 17 CONSECUTIVE 10 SECOND TIME SERIES. INDEX 05 CORRESPONDS TO THE DATA USED IN TABLE I. THE TABLE SHOWS THE ESTIMATED PARAMETERS AND 1 STANDARD DEVIATION OF THAT ESTIMATE (SD) FOR THE WIENER AND AR(p) MODELS. THE BOTTOM ROW SHOWS THE AVERAGE ESTIMATE, WEIGHTED BY THE VARIANCE OF EACH INDIVIDUAL ESTIMATE (DENOTED BY WE FOR WEIGHTED ESTIMATE).

	Wiener		**AR(p)**											
	σ^2	SD	θ_1	SD	ϕ_1	SD	ϕ_2	SD	ϕ_3	SD	ϕ_4	SD	ϕ_5	SD
01	.7057	.0075	-.9073	.0044	.4348	.0084	.1546	.0077	-	-	-	-	-	-
02	.8513	.0085	-.8827	.0046	.4891	.0086	-	-	-	-	-	-	-	-
03	.8587	.0086	-.8960	.0042	.5259	.0083	-	-	-	-	-	-	-	-
04	.6365	.0071	-.9175	.0039	.5591	.0081	.0610	.0075	-	-	-	-	-	-
05	.5712	.0061	-.9434	.0026	.5605	.0075	.0428	.0073	-	-	-	-	-	-
06	.6548	.0071	-.8779	.0054	.4967	.0089	.0302	.0079	-	-	-	-	-	-
07	.8400	.0085	-.8793	.0048	.4764	.0088	-	-	-	-	-	-	-	-
08	.8422	.0084	-.8928	.0044	.5562	.0084	-	-	-	-	-	-	-	-
09	.5457	.0059	-.9452	.0028	.4806	.0075	.1648	.0078	.0064	.0073	-	-	-	-
10	.5354	.0058	-.9334	.0033	.4889	.0078	.1477	.0075	-	-	-	-	-	-
11	.7301	.0076	-.9291	.0046	.5100	.0085	.0194	.0080	-.0413	.0080	.1836	.0075	-	-
12	.7280	.0076	-.9156	.0056	.4896	.0090	.0334	.0082	.0025	.0081	.0942	.0080	.0024	.0074
13	.7885	.0081	-.9191	.0051	.4175	.0087	.1326	.0080	-.0171	.0078	.0914	.0076	-	-
14	.8875	.0095	-.9036	.0092	.5500	.0114	.0318	.0086	.0447	.0084	.1329	.0082	.0308	.0080
15	.6392	.0076	-.9292	.0047	.5535	.0084	.1231	.0082	.0557	.0081	.0067	.0073	-	-
16	.6587	.0076	-.9357	.0049	.5649	.0086	.0960	.0083	.0570	.0082	.0298	.0074	-	-
17	.6767	.0077	-.9075	.0050	.6108	.0086	.0465	.0083	.0029	.0075	-	-	-	-
WE	.6859		-.9200		.5153		.0845		.0126		.0874		.0155	

B. Wiener Process

We can fit a Wiener process model to the data as well. As the data already represents the first difference of the phase noise process, we can use the standard unbiased variance estimator to estimate σ^2:

$$\hat{\sigma}^2 = \frac{1}{N-1} \sum_{k=1}^{N} \phi_{\Delta}[k] \tag{5}$$

For the data used in Table I, we estimate that it is best modeled as coming from a Wiener process parameterized by $\hat{\sigma}^2 = 0.5712$ per $10\mu s$.

C. Consistency of Models

The preceding analysis has been performed on a single data set representing 10 seconds. As we do not expect the phase noise process to be necessarily stationary over longer time, we must investigate how these models hold over time. For the Wiener process model, we must determine if the σ^2 value holds over time or if it must be re-estimated frequently. The AR model raises even more questions:

- Does the order of the process have to be re-estimated?
- Do the process parameters have to be re-estimated?

We begin to address these concerns by examining the short time stationarity of the phase noise. Table II show analysis results of 17 consecutive 10 second analyses fitted to Wiener and AR(p) models. We see that the estimate of σ^2 does appear to drift a bit but generally stays close to the weighted average value. In the case of the AR model, we see that the θ_1 estimate of the differenced time series also stays around the weighted average value. The order, p, varies from 1 to 5, but the estimates of ϕ_1 appear to roughly agree with each other across time. The values of the higher order parameters, when they exist, decrease in magnitude as p increases. This seems to indicate that they are less important and possibly show up only due to variability in additional noise sources, suggesting that an AR(1) model is appropriate across this range of time. Even a single estimate of ϕ_1 may suffice (instead of re-estimating that parameter for every 10 second data set).

Table III shows the same short time evolution of parameter estimates for Wiener and AR(p) models but for data recorded approximately 2 years earlier than the data in Table II. This allows us to examine the longer term stationarity of the phase noise process. We see that the estimate of σ^2 for the Wiener model is noticeably lower. The AR modeling shows that a $p = 1$ order still appears to hold, and even the estimates of ϕ_1 are on par with the estimates from Table II. The estimates of θ_1 for this data are closer to -1, indicating that this data is better represented by a pure AR process than the data from Table II.

IV. CONCLUSION

We have shown that both AR and Wiener process models appear to model the phase noise data from MRR oscillators well. The Wiener process model results indicated that should such a model be used, the σ^2 parameter should be estimated adaptively, updating the estimate as needed as it seems to fluctuate even over short times. The AR model results suggest that an order $p = 1$ process works well even over long times, and although estimates of the AR parameters appear to hold over long time, the model would be well served to have the AR parameter estimates updated adaptively as well.

TABLE III

TIME SERIES ANALYSIS RESULTS FOR 17 CONSECUTIVE 10 SECOND TIME SERIES RECORDED APPROXIMATELY 2 YEARS PRIOR TO THE DATA IN TABLE II. THE TABLE SHOWS THE ESTIMATED PARAMETERS AND 1 STANDARD DEVIATION OF THAT ESTIMATE (SD) FOR THE WIENER AND AR(p) MODELS. THE BOTTOM ROW SHOWS THE AVERAGE ESTIMATE, WEIGHTED BY THE VARIANCE OF EACH INDIVIDUAL ESTIMATE (DENOTED BY WE FOR WEIGHTED ESTIMATE).

	Wiener		**AR(p)**									
	σ^2	SD	θ_1	SD	ϕ_1	SD	ϕ_2	SD	ϕ_3	SD	ϕ_4	SD
01	.5041	.0050	-.9977	.0002	.5030	.0060	-	-	-	-	-	-
02	.4215	.0042	-.9986	.0007	.4272	.0070	.1417	.0070	-	-	-	-
03	.4268	.0043	-.9982	.0003	.5593	.0058	-	-	-	-	-	-
04	.5151	.0050	-.9976	.0002	.3563	.0069	.1839	.0069	-	-	-	-
05	.5310	.0055	-.9742	.0015	.3486	.0071	.0929	.0074	.1245	.0071	-	-
06	.3454	.0048	-.9962	.0004	.6143	.0070	-.1518	.0082	.1204	.0082	-.0056	.0070
07	.3565	.0047	-.9717	.0015	.5640	.0061	-	-	-	-	-	-
08	.3805	.0048	-.9688	.0016	.6473	.0058	-	-	-	-	-	-
09	.5465	.0062	-.9624	.0017	.5277	.0065	-	-	-	-	-	-
10	.5897	.0064	-.9624	.0017	.5040	.0065	-	-	-	-	-	-
11	.5411	.0060	-.9966	.0003	.5081	.0060	-	-	-	-	-	-
12	.5659	.0061	-.9977	.0002	.5079	.0060	-	-	-	-	-	-
13	.5400	.0060	-.9506	.0021	.5532	.0066	-	-	-	-	-	-
14	.5565	.0061	-.9976	.0003	.4820	.0061	-	-	-	-	-	-
15	.3548	.0045	-.9985	.0005	.6669	.0070	-.0510	.0070	-	-	-	-
16	.4114	.0047	-.9787	.0012	.6337	.0057	-	-	-	-	-	-
17	.3752	.0044	-.9991	.0012	.6366	.0062	-	-	-	-	-	-
WE	.4491		-.9969		.5371		.0543		.1227		-.0056	

The estimates of ϕ_1 seem counter-intuitive at first. With high quality oscillators, we expected an AR(1) model to have a larger estimate (something more like 0.9 instead of the values around 0.5 that we are seeing). Further analysis reveals that this may be due to allowing the ARMA model to use θ_1 values slightly larger than -1. Once we had determined that the appropriate AR order to use is 1, we attempted to estimate ϕ_1 by fitting the known spectral density function for an ARMA(1,1) model, setting $\theta_1 = -1$, and allowing ϕ_1 to fluctuate while minimizing an error function. This procedure results with estimates of ϕ_1 in the expected neighborhood of 0.9.

We can follow the estimates of these parameters to judge system performance. In the AR model, the closer $\hat{\phi}_1$ is to -1, the better the system is performing (or at least, the less affected by phase noise it is). In the Wiener model, smaller values of $\hat{\sigma}^2$ correspond to less phase noise.

These stochastic models can also be used in parametric approaches of phase noise correction. Autofocus algorithms have already been developed that deal with phase noise correction in atmospheric distortion in terrestrial telescope imaging [8] and platform motion compensation in synthetic-aperture radar (SAR) [2]. These techniques work by developing an image sharpness metric and then maximizing that metric [4]. Both parametric and non-parametric approaches to autofocus algorithms exist. As oscillator phase noise is relatively small compared to phase noise in the previous examples, parametric approaches based on the stochastic models we have presented are advantageous in reducing the complexity of the optimization routines especially since the phase noise itself, if small enough, may approach the scale of the convergence noise of the optimization routine.

REFERENCES

[1] A. Behravan, F. Munier, T. Svensson, M. Flament, T. Eriksson, A. Svensson, and H. Zirath. System implications in designing a 60 GHz WLAN RF front-end. In *Proc. of GigaHertz 2001 Symposium on Gigahertz Electronics*, Nov. 2001.

[2] T.M. Calloway and G.W. Donohoe. Subaperture autofocus for synthetic aperture radar. *IEEE Trans. on Aero. and Elec. Sys.*, 30(2), Apr. 1994.

[3] L.S. Cutler and C.L. Searle. Some aspects of the theory and measurement of frequency fluctuations in frequency standards. In *Proc. of the IEEE*, Feb. 1966.

[4] J.R. Fienup and J.J. Miller. Aberration correction by maximizing generalized sharpness metrics. *J. Opt. Soc. Am.*, 20(4), Apr. 2003.

[5] A. Hajimiri and T.H. Lee. A general theory of phase noise in electrical oscillators. *IEEE Journal of Solid-State Circuits*, 33(2), Feb. 1998.

[6] V.F. Kroupa. Noise properties of PLL systems. *IEEE Trans. on Comm.*, Com-30(10), Oct. 1982.

[7] M.G. Meyer, J.D. Sahr, and A.N. Morabito. A statistical study of subauroral E-region coherent backscatter observed near 100 MHz with passive radar. *JGR*, 109, doi:10.1029/2004JA010396, 2004.

[8] R.A. Muller and A. Buffington. Real-time correction of atmospherically degraded telescope images through image sharpening. *J. Opt. Soc. Am.*, 64(9), Sep. 1974.

[9] M. Okumura and H. Tanimoto. A time-domain method for numerical noise analysis of oscillators. *Proc. ASP-DAC*, Jan. 1997.

[10] T. Pollet, M. Van Bladel, and M. Moeneclaey. BER sensitivity of OFDM systems to carrier frequency offset and Wiener phase noise. *IEEE Trans. Comm.*, 43(2/3/4), Feb/Mar/Apr 1995.

[11] V.C. Vannicola. Spectral dispersion of modulated signals due to oscillator phase instability: White and random walk phase model. *IEEE Trans. on Comm.*, Com-31(7), Jul. 1983.