Signal Acquisition and Tracking of Chirp-Style GPS Jammers

Ryan H. Mitch, Mark L. Psiaki, Steven P. Powell and Brady W. O'Hanlon,
Cornell University, Ithaca, NY

BIOGRAPHIES

Ryan H. Mitch is a Ph.D. candidate in the Sibley School of Mechanical and Aerospace Engineering at Cornell University. He received his B.S. in Mechanical Engineering from the University of Pittsburgh and his M.S. in Mechanical Engineering from Cornell University. His current research interests are in the areas of GNSS technologies and integrity, nonlinear estimation and filtering, and signal processing.

Mark L. Psiaki is a Professor in the Sibley School of Mechanical and Aerospace Engineering. He received a B.A. in Physics and M.A. and Ph.D. degrees in Mechanical and Aerospace Engineering from Princeton University. His research interests are in the areas of GNSS technology, applications, and integrity, spacecraft attitude and orbit determination, and general estimation, filtering, and detection.

Steven P. Powell is a Senior Engineer with the GPS and Ionospheric Studies Research Group in the Department of Electrical and Computer Engineering at Cornell University. He has M.S. and B.S. degrees in Electrical Engineering from Cornell University. He has been involved with the design, fabrication, testing, and launch activities of many scientific experiments that have flown on high altitude balloons, sounding rockets, and small satellites. He has designed ground-based and space-based custom GPS receiving systems primarily for scientific applications.

Brady W. O'Hanlon is a Ph.D. candidate in the School of Electrical and Computer Engineering at Cornell University. He received both his M.S. and B.S. in Electrical and Computer Engineering from Cornell University. His interests are in the areas of GNSS technology and applications, GNSS security, and GNSS as a tool for space weather research.

ABSTRACT

This paper investigates one method of acquiring and one method of tracking various chirp-style GPS jammers. A signal model that uses polynomial descriptions of frequency versus time is developed. The developed model can track ideal linear chirps as well as more complicated, and even some non-repeating, polynomial frequency patterns. The developed model can also be used for jammer classification. Two slightly different measurement models that both make use of Fast Fourier Transforms (FFTs) are also developed. An ad-hoc chirp-style signal acquisition method that uses FFTs of a moderately strong jamming signal is also developed. The jamming signal model, observation model, and acquisition procedure are combined to create a chirp-style signal tracking Kalman filter. The developed Kalman filter is verified by application on three sets of laboratory data and on two sets of field data. Tracking of a chirp-style jammer is demonstrated for a distance of approximately 1.8 kilometers between the receiving station and the jammer. The jammer signal models and the Kalman filter are also expanded to the scenario where multiple jamming signals are present, and the modified Kalman filter is evaluated on one set of laboratory data.

INTRODUCTION

The Global Positioning System's location and time-synchronization capabilities are used in many areas of civilian life. Some common civilian uses include navigation through GPS-enabled smart phones or dash-mounted navigation units, and geotagging photographs. Common commercial uses include tracking trucking and shipping [1], aircraft and maritime navigation [2], and high precision timing applications [3, 4].

Preprint from ION GNSS+ 2013

Government agencies, such as the police or FBI, can use GPS for tracking suspected criminals [5].

Unfortunately, the interests of some individuals can be served by interfering with GPS. A simple example is that of a thief who steals a vehicle that is GPS enabled and wishes to interfere with GPS so that the vehicle cannot be recovered before dismantling. Another example is that of an employee in a commercial trucking corporation who wishes to interfere with the GPS tracking device on his company truck so that he may run personal errands while being paid to make deliveries. A less malignant use of GPS interference would be that of someone attempting to enforce an envelope of privacy around their personal vehicle [6].

In the above examples the GPS interference could be provided by a civil GPS jammer, also known as a Personal Privacy Device (PPD). This has led to several incidents, of which the so called "Newark Incident" is the most commonly recognized. In the Newark Incident a truck driver with a GPS jammer in his vehicle drove by Newark airport and periodically interfered with the airport's GPS equipment [2]. The truck driver only wanted to enforce an envelope of privacy around his vehicle and was not intending to interfere with the airport equipment. There was also a less benign incident in Great Britain where a group of car thieves used GPS jammers to try to disrupt the geolocation and recovery efforts of the relevant authorities [1].

The above incidents have motivated a number of researchers to investigate PPDs [7, 8, 9, 10, 11, 12, 13] in general and their geolocation [14, 15, 16, 17, 18] in specific. This paper furthers the work of [15] and provides a set of algorithms to acquire and track various chirp-style GPS jammers. Jammer signal tracking has applications in jammer geolocation [15], as would be useful for law-enforcement actions. Although, not investigated in detail, this work may also have applications in jammer classification and interference mitigation.

The remainder of the paper is divided into eight sections. The first section presents background information on the civilian GPS chirp-style jammers. The second section develops a model and state parameterization for the PPDs' chirp-style signals. The third section briefly discusses jammer classification using the developed models. The fourth section discusses and selects an observation model for use in jammer signal tracking. The fifth section outlines an ad-hoc strategy for acquiring a chirp-style jammer that has a moderately strong carrier-to-noise ratio. The sixth section combines the state parameterization, observation model, and acquisition procedure and applies the resulting algorithm to data collected from several individual jammers. The seventh section extends the jammer modeling to scenarios where multiple jammers are present and it also considers the new complications that arise when multiple jamming signals must be tracked. Additionally, the section applies the multi-jammer tracking algorithm to multi-jammer laboratory data. The final section summarizes the paper's developments and draws appropriate conclusions.

GPS JAMMER BACKGROUND INFORMATION

Civilian GPS jammers/PPDs can be found in a variety of form factors, but are on average approximately the size of a hand-held cellular telephone [8]. Three different civilian GPS jammers are shown in Fig. 1.

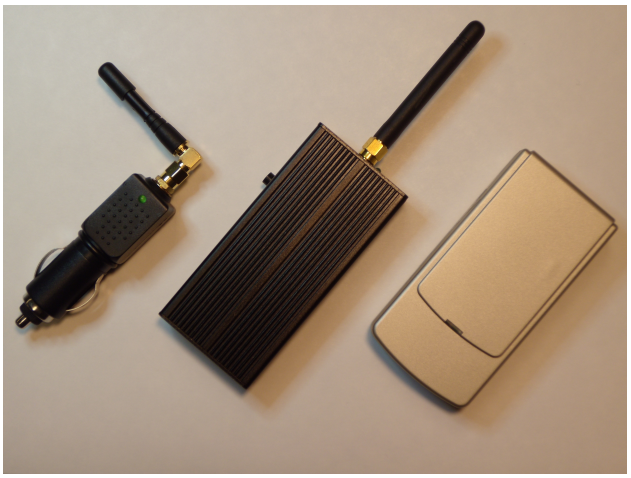

Figure 1 Three different form factors of civilian GPS jammers/PPDs.

The algorithms used in the processing of signals from GPS jammers can benefit from an understanding of the RF output of those same jammers. The typical output of a civil GPS jammer is shown in Fig. 2. The horizontal axis is time and the top plot's vertical axis is frequency. Each vertical slice of the top plot in the figure is a Fast Fourier Transform (FFT) of the RF sampled signal, centered at the GPS L_1 frequency. The z, or color axis, is power, with red denoting a large value and blue denoting a small value. The bottom plot's vertical axis is power. The figure shows a classic example of a chirp signal, or a tone whose frequency repeatedly ramps linearly upwards and then resets back to the starting frequency.

The plot shown in Fig. 2 is a common example of a GPS jammer's output spectrum, but other minor variations exist and will be addressed later. Further infor-

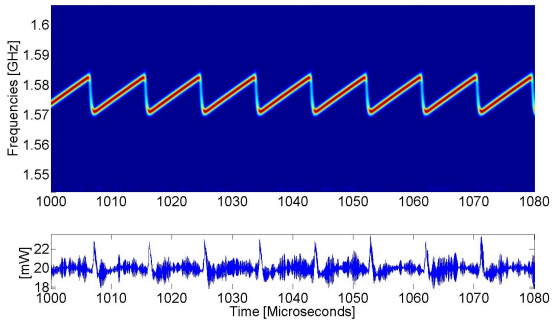

Figure 2 A common GPS jammer spectrum. The top plot displays vertical slices of 64-point Hamming-windowed FFTs, and the bottom plot is of power.

mation can be found in the survey of civilian GPS jammers in Ref. [8].

JAMMER MODELING AND PARAMETERIZATION

There are multiple parameterizations that could be developed for the chirp-style jammer signal introduced in the previous section. This paper considers a parameterization that is similar to that used in Ref. [15], but with several modifications. The parameterization starts by assuming a linear chirp, or a first-order rate of change of the frequency versus time, as shown in Fig. 3.

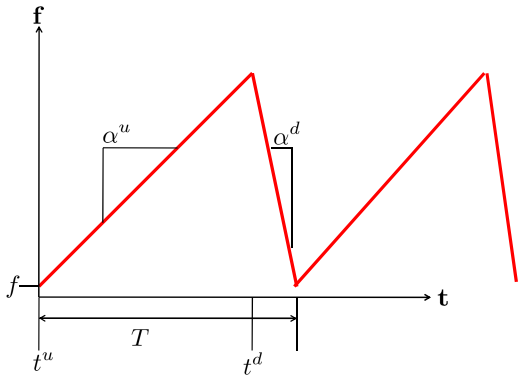

Figure 3 Time-history of the signal frequency of a linear first-order chirp, with appropriate states labeled.

The following state vector is a moderately low-order parameterization of a chirp-style jammer:

$$\underline{x} = \begin{bmatrix} \theta \\ f \\ \alpha^u \\ \alpha^d \\ A \\ t^u \\ t^d \\ T \end{bmatrix} \tag{1}$$

where the entries of the state vector are as follows: θ is the phase in units of cycles, f is the frequency in units of Hertz, α^u is the upward frequency rate of change in units of Hertz per second, α^d is the downward frequency rate of change in units of Hertz per second, A is the amplitude in units of Volts, t^u is the ramp up time start for the current ramp and is in units of seconds, t^d is the ramp down time start for the current ramp and is in units of seconds, T is the chirp period and is in units of seconds. The ramp times could have separate periods, such as T^u and T^d, but experimentation with separate periods did not dramatically change the results presented later in this paper.

The above parameterization assumes that the jammer chirp has a linear first-order polynomial rate of change of frequency versus time. For the majority of the chirp-style jammers the above parameterization is a sufficiently accurate model for most signal processing algorithms. However, there are several jammers that have significantly different behavior. The following two figures show the spectra of two jammers that are not ideally approximated by linear chirps. Fig. 4 appears to have a non-first-order polynomial ramp down in frequency, whereas Fig. 5 has a non-first-order ramp up in frequency.

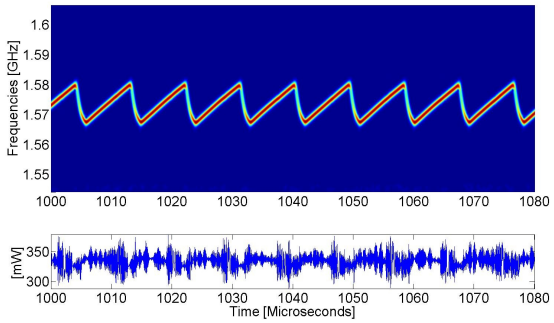

Figure 4 A jammer power spectra, with a non–first-order polynomial ramp down in frequency versus time.

The model state can be modified in several ways to account for higher-order variations in the frequency

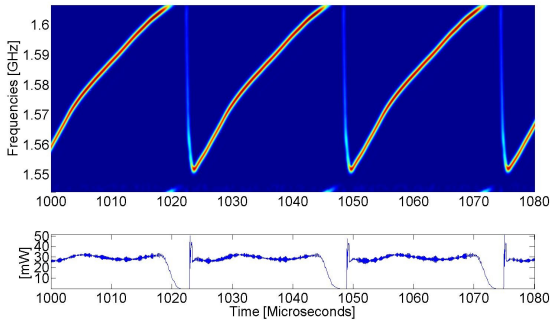

Figure 5 A jammer power spectra, with a non–first-order polynomial ramp up in frequency versus time.

ramps. The modification method selected in this paper is to add new states that allow the frequency ramps to follow higher-order polynomials, as is shown in Fig. 6. The polynomial expansion is in some ways analogous to a Taylor series expansion of a nonlinear term.

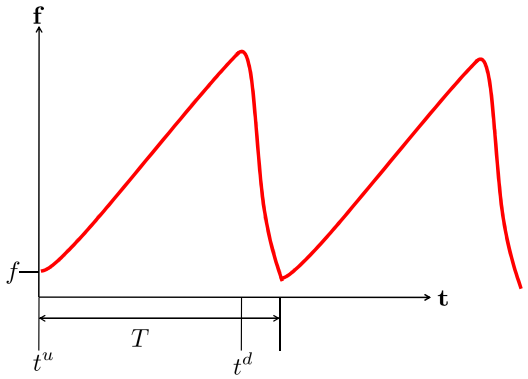

Figure 6 Possible time-history of the signal frequency for a polynomial-type chirp, with appropriate states labeled.

The general form of the new state vector is as follows:

$$\underline{x} = \begin{bmatrix} \theta \\ f \\ \begin{bmatrix} c_1^u \\ \vdots \\ c_{M_u}^u \end{bmatrix} \\ \begin{bmatrix} c_1^d \\ \vdots \\ c_{M_d}^d \end{bmatrix} \\ A \\ t^u \\ t^d \\ T \end{bmatrix} \tag{2}$$

All of the c_j^u and c_j^d states are coefficients of the Taylor-series-type polynomial expansion of the frequency behavior, including the first order terms c_1^u and c_1^d that have replaced α^u and α^d. The states of the form c_j^u are the coefficients of a ramp-up frequency polynomial, and the states of the form c_j^d are the coefficients of a ramp-down frequency polynomial. The frequency polynomial is defined as follows:

$$f\left(\underline{x},t\right) = \begin{cases} f + \sum_{j=1}^{M_u} c_j^u \left(t - t^u\right)^j, & t_u \leq t_d \\ f + \sum_{j=1}^{M_d} c_j^d \left(t - t^d\right)^j, & t_d < t_u \end{cases} \tag{3}$$

and it can be integrated to determine the corresponding phase:

$$\theta\left(\underline{x},t\right) = \begin{cases} \theta + f\left(t - t^u\right) + \sum_{j=1}^{M_u} c_j^u \frac{\left(t-t^u\right)^{(j+1)}}{(j+1)}, & t_u \leq t_d \\ \theta + f\left(t - t^d\right) + \sum_{j=1}^{M_d} c_j^d \frac{\left(t-t^d\right)^{(j+1)}}{(j+1)}, & t_u > t_d \end{cases} \tag{4}$$

where the top case in Eqs. 3 and 4 is for times when the frequency is ramping upwards and the bottom case is for times when the frequency is ramping downwards.

The order of the polynomial, and the associated states c_1^u–$c_{M_u}^u$ and c_1^d–$c_{M_d}^d$, can be considered a tuning parameter. More states will allow for a model that can more closely replicate the behavior of the true jammer. The practical upper bound on the number of coefficients used in the model is primarily set by the sampling rate and numerical conditioning limitations. If the sampling rate is very low, then the number of accumulations that can be computed for each chirp is

reduced and it will be difficult to estimate the higher order terms accurately. Numerical conditioning issues may arise due to the dramatically different state magnitudes in the above parameterization, and this may require that the states be modified to improve the numerical stability of the algorithm. It should also be noted that using more polynomial coefficients than necessary is a waste of computational effort. Therefore, the order of the polynomial should be limited to a reasonable number.

A 5^{th} order polynomial is considered for both the frequency ramp-up and ramp-down in the remainder of this paper's modeling work. Additionally, the polynomial coefficients have been modified from the definition in Eqs. 2–4. The new polynomial coefficients have been multiplied by the sample time T_s, raised to the corresponding power, as shown below for the ramp up:

$$\begin{aligned} c^u_{1,new} &= c^u_1 (T_s)^1 \\ &\vdots \\ c^u_{M_u,new} &= c^u_{M_u} (T_s)^{M_u} \end{aligned} \tag{5}$$

and for the ramp down:

$$\begin{aligned} c^d_{1,new} &= c^d_1 (T_s)^1 \\ &\vdots \\ c^d_{M_d,new} &= c^d_{M_d} (T_s)^{M_d} \end{aligned} \tag{6}$$

The new states will have magnitudes that are closer to unity, and therefore will be less likely to cause numerical instabilities in the developed algorithms. For example, a reasonable numerical value for c^u_1 might be 10^{12}, and a system with a sampling rate of 10 MHz would cause $c^u_{1,new}$ to have a numerical value closer to 10^5. It would also be possible to use a fixed number that is unrelated to sample rate.

The state dynamics are defined in a chirp-by-chirp manner, with the frequency and phase defined at the ramp times t^u and t^d. The phase and frequency can be evaluated at any time, but are only propagated forward and updated when the time crosses the ramp-down time, t^d, or ramp-up time, t^u, by evaluating Eqs. 3 and 4.

The resulting state dynamics equation is as follows:

$$\underline{x}_{k+1} = \Phi_k \left(t_{k+1}, t_k; \underline{x}_k\right) \underline{x}_k + \Gamma_k \left(t_{k+1}, t_k; \underline{x}_k\right) \underline{v}_k \tag{7}$$

where $\underline{x}_{k+1}$ is the state at time t_{k+1}, $\underline{x}_k$ is the state at time t_k, Φ_k is the state transition matrix from time t_k to t_{k+1}, Γ_k is the process noise influence matrix, and $\underline{v}_k$ is the zero-mean, unity-covariance, white, Gaussian process noise vector. The process noise is expected to enter only at the beginning of each ramp interval, t^u or t^d. It should be noted that the state transition matrix and the process noise influence matrix are both state dependent, and are defined as follows:

$$\Phi \left(t_{k+1}, t_k; \underline{x}_k\right) = \begin{cases} \Phi \left(t_{k+1}, t_k\right)_{\text{up}}, & \text{if A} \\ \Phi \left(t_{k+1}, t_k\right)_{\text{down}}, & \text{if B} \\ \text{I}, & \text{otherwise} \end{cases} \tag{8}$$

$$\Gamma \left(t_{k+1}, t_k; \underline{x}_k\right) = \begin{cases} \Gamma \left(t_{k+1}, t_k\right)_{\text{up}}, & \text{if A} \\ \Gamma \left(t_{k+1}, t_k\right)_{\text{down}}, & \text{if B} \\ \underline{\underline{0}}, & \text{otherwise} \end{cases} \tag{9}$$

where the conditions A and B are defined as follows:

$$\text{Condition A: } t = t^d \tag{10}$$

$$\text{Condition B: } t = t^u \tag{11}$$

The state transition matrices for this system are the identity matrices, except when Eqs. 10 and 11 are satisfied, i.e. when the system time equals t^d or t^u. When the system time equals t^d the phase and frequency are propagated as defined by the up-ramp polynomial from time t^u up to time t^d. Similarly, when the system time equals t^u the phase and frequency are propagated as defined by the down-ramp polynomial from time t^d up to time t^u. The time difference terms in Eqs. 3 and 4 can be rewritten as rows of $\Phi \left(t_{k+1}, t_k\right)_{\text{up}}$ or $\Phi \left(t_{k+1}, t_k\right)_{\text{down}}$ because the states, the coefficients of the polynomials, enter linearly into those equations. The amplitude, A, is defined as a constant and the time states are defined by the following equations:

$$t_u = \begin{cases} t_u + T, & \text{if } t = t_d \\ t_u, & \text{otherwise} \end{cases} \tag{12}$$

$$t_d = \begin{cases} t_d + T, & \text{if } t = t_u \\ t_d, & \text{otherwise} \end{cases} \tag{13}$$

where the time states t^u and t^d are incremented by their period T when the system time equals t^d or t^u, respectively. The rows corresponding to t^u or t^d in $\Phi \left(t_{k+1}, t_k\right)_{\text{up}}$ or $\Phi \left(t_{k+1}, t_k\right)_{\text{down}}$ will have 1 entries for that state, and if the conditions of Eqs. 10 or 11 are satisfied, it will also have a 1 entry for the period increment for the corresponding state.

Propagation of the state from a time before t^u or t^d to a time after t^u or t^d will require multiple matrices. For example, if the state were propagated from time t_k to t^u and then from t^u to t_{k+1} the resulting equations would be:

$$\begin{aligned} \Phi \left(t_{k+1}, t_k, \underline{x}_k\right) &= \Phi \left(t_{k+1}, t^u_+\right) \Phi \left(t^u_+, t^u_-\right) \Phi \left(t^u_-, t_k\right) \\ &= \text{I } \Phi \left(t_{k+1}, t_k\right)_{\text{down}} \text{ I} \\ &= \Phi \left(t_{k+1}, t_k\right)_{\text{down}} \end{aligned} \tag{14}$$

where the time t^u_- designates a time infinitesimally before t^u and time t^u_+ designates a time infinitesimally after t^u. If multiple ramp changes were included in the propagation interval then matrix multiplication series of Eq. 14 would contain more terms.

The default process noise influence matrices for this system are zero matrices, except when the conditions in Eqs. 10 or 11 are met, i.e. when the system time equals t^d or t^u. All of the states, except for θ, are considered to have some process noise enter at the beginning of each ramp. The introduction of process noise creates non-zero entries in the Γ matrices, which are assumed to be diagonal. If the non-zero diagonal entries are each set to 1, then the process noise covariance matrix Q becomes the square of the standard deviation of the process noise expected for the entire ramp. Q is formally defined as:

$$Q = E[(\underline{v}_k - E[\underline{v}_k])\,(\underline{v}_k - E[\underline{v}_k])^*] = E[\underline{v}_k \underline{v}_k^*] \quad (15)$$

where $E\,[]$ is the expected value of the quantity in the brackets. A reasonable assumption for the Q matrix is that all of the noise inputs are uncoupled, i.e. Q is a diagonal matrix.

The amount of process noise that is assumed to be present in each signal state is a tuning parameter that can be modified to improve the performance of the filter. The signal tracking results of the single GPS jammers were not extremely sensitive to the assumed level of process noise. However, the multi-jammer scenarios were decidedly more sensitive to the process noise tuning. It should be noted that each jammer can have a different level of process noise, and a thorough survey could be conducted to determine a statistically significant estimate of the process noises that could be expected from any GPS jammer seen in the field. A survey of that many jammers is beyond the scope of this current work.

The resulting output of the above jammer model is:

$$y'_i = A'_i * cos(2\pi * \phi(\underline{x}, t_i)) \quad (16)$$

where $\phi(\underline{x}, t_i)$ is the evaluation of the relevant phase polynomial, and A'_i is the broadcast signal amplitude, at time t_i. The signal as seen at a receiver station with mixing frequency f_{mix}, and with the high-frequency mixing term removed, is as follows:

$$\begin{aligned} y_i &= \frac{A_i}{2} * cos\left(2\pi * \left(\phi\left(\underline{x}, t_i\right) - f_{mix} t_i\right)\right) \\ &= \frac{A_i}{2} * cos\left(\Theta\left(\underline{x}, t_i, f_{mix}\right)\right) \end{aligned} \quad (17)$$

where the amplitude in Eq. 17 is the actual tracked amplitude state. Eq. 17 is the final form of the signal model that is used in the remainder of this paper.

JAMMER CLASSIFICATION

In Ref. [8] the authors found that jammers with the same physical appearances produced significantly different frequency versus time behavior. The unique frequency versus time behavior can potentially be used to identify individual jammers. Jammer identification can be useful for determining how many different jammers are seen at one station. It can also be useful if a person is being prosecuted for operating a PPD and the authorities wish to know where else that person has broken the law by jamming GPS. Jammer identification only requires that at least one station receive the jamming signal, as opposed to most automated geolocation methods which require several stations.

One possible identification method would be to store a short amount of time of high-sample-rate data from many jammers in an RF signature library and then compute cross correlations between the received data and the library to determine if the signals match. A new method using the modeling in this paper would also start with a short recording of data from each jammer in the library. The jammer state would then be estimated for many chirps. The resulting time-series of state estimates would then comprise many samples from a probability density function (pdf) of the jammer's state perturbed by process noise. A similar procedure would then be applied to the new received signal and the resulting two sets of pdfs would be compared to each other. Relevant similarity metrics and hypothesis tests could be developed to determine if the new signal is the same as one of the signals in the library. This method of jammer identification might require a smaller amount of storage space when compared to one which used a sampled RF data library. This is a potentially deep research area and will require further investigation to explore.

JAMMER OBSERVABLES

The standard pair of observables used to track an RF signal are the In-Phase and Quadrature accumulations. The accumulation formulas at time t_k are:

$$I_k = \sum_{i=1}^{N} y_i cos\left(\Theta^{NCO}\left(t_i\right)\right) \quad (18)$$

$$Q_k = \sum_{i=1}^{N} y_i sin\left(\Theta^{NCO}\left(t_i\right)\right) \quad (19)$$

where N is the number of samples used in the accumulation, y_i is the sampled RF data, and $\Theta^{NCO}\left(t_i\right)$ is

the numerically controlled oscillator (NCO) phase at time t_i.

The NCO phase is arbitrarily defined by the user. The jammer phase time history can be determined from the previous modeling section and is as follows for a single polynomial ramp:

$$\Theta\left(\underline{x}, t_i, f_{mix}\right) = 2\pi\left[\theta_k^N + f_k^N \Delta t_A + \left(\sum_{j=1}^{N_A} c_j^A \frac{\Delta t_A^{j+1}}{(j+1)\left(T_s\right)^j}\right) - f_{mix} t_i\right] \tag{20}$$

where the subscript "A" denotes ramp A, Δt_A is defined as $t_i - t_A$, and t_A is either t^u or t^d, depending on if the ramp is up or down, respectively. If the accumulation time interval spans part of two ramps, ramp A and then ramp B, the phase is as follows:

$$\Theta\left(\underline{x}, t_i, f_{mix}\right) = 2\pi\left[\theta_k^N + f_k^N \Delta t_A + \left(\sum_{j=1}^{N_A} c_j^A \frac{\Delta t_{BA}^{j+1}}{(T_s)^j}\left[\frac{\Delta t_{BA}}{(j+1)} + \Delta t_B\right]\right) + \left(\sum_{j=1}^{N_B} c_j^B \frac{\Delta t_B^{j+1}}{(j+1)\left(T_s\right)^j}\right) - f_{mix} t_i\right] \tag{21}$$

where Δt_{BA} is defined as $t_B - t_A$, and t_B is either t^u or t^d and t_A is the other, i.e. if the first ramp is an up ramp then ramp A starts at time t^u and $t_B = t^d$. If more than two ramps are spanned by an accumulation interval then the received phase equation simply gains additional terms.

If the signal contains no noise, then the NCO phase that produces the largest accumulation is the same phase as the incoming signal that one wishes to track. In a traditional signal tracking architecture the accumulations would then constitute the measurements that would be sent to the Kalman filter. The corresponding measurement models would then be defined by substituting the signal model and the signal model estimates into the y_i term in the accumulation models of Eqs. 18 and 19. The measurement models can be simplified further by converting the discrete sums to continuous integrals and attempting to solve the resulting integrals in closed form. The integral of the higher order phase terms is at simplest, in the case of a perfectly linear ramp, a Fresnel integral. The higher order polynomials will further complicate the integration. The resulting measurement model is complicated and requires significant computational overhead. Additionally, a single pair of long time-span accumulations is likely to prove insufficient for filter convergence and signal tracking. This is because the accumulations will result in appreciable power only if the NCO phase time-histories are close to that of the real signal being tracked.

The slow evaluation speed and a strong preference for multiple accumulations suggest that a different measurement model that does not have the same drawbacks should be considered. It should be noted that, unlike GPS, the GPS jammers broadcast very powerful signals, and their signal can be easily seen above the noise floor in a given area around each jammer. A faster measurement that computes accumulations and spans the frequency range of the physical sampling system is the Fast Fourier Transform (FFT). The FFT computes accumulations at fixed frequencies spaced evenly between the positive and negative Nyquist frequency bounds. If the accumulation lengths are short enough that they do not lose power due to the changing frequency of the incoming signal then they will provide a measurement of the jammer power in each frequency bin across the entire Nyquist range.

The FFT measurements are defined as follows:

$$\begin{aligned} z(L) &= I + iQ \\ &= \sum_{n=1}^{N} w(n) y(n) e^{-i2\pi \frac{(L-1)(n-1)}{N}} \\ &= \sum_{n=1}^{N} w(n) y(n) e^{-i\zeta(L,n,N)} \\ &\quad \forall L \in [1, ..., N] \end{aligned} \tag{22}$$

where w is the time domain windowing function used to reduce sidelobes in the FFT, y is the sampled signal, and L is the frequency bin number of the FFT. The present work used the Hamming window, but other windowing functions would likely produce similar results.

It is important to note that it would be unreasonable to expect that the phase can be precisely tracked in this system. Precise tracking of the phase state θ requires at least the following two conditions be met. The first condition is that the frequency polynomial estimate be very accurate, such that there are only very minor fluctuations in the frequency that cannot be modeled by the polynomial. The second condition is that the jamming signal always stays within the receiver's Nyquist frequency bounds.

Satisfying the above two conditions will be difficult in even the most benign situations, where the jamming environment can be controlled and RF data can be acquired with a high-speed recording system. The results of two such scenario are shown in Figs. 7 and 8. In Fig. 7 the order of the frequency polynomial required to describe the signal changes from chirp to

chirp, and sometimes the frequency of the signal performs near-instantaneous jumps that may not be accurately modeled by a continuous polynomial. Even a well-designed signal tracking algorithm will therefore be unable to accurately track this jammer's phase, unless the instantaneous jump behavior is accurately modeled and reliably estimated. The frequency polynomial model developed in this paper does not permit instantaneous frequency jumps, but it will allow an approximation to that behavior. In Fig. 8 the frequency of the jamming signal passes outside of the receiver system's Nyquist range. When the frequency passes outside of the Nyquist range the phase can no longer be tracked. Although the frequency and phase behavior are governed by the polynomial parameterizations developed in this paper, they are not guarantees of the behavior outside of the visible frequency range. The Nyquist frequency range is dependent on the equipment being used, but even in the above high-rate laboratory system the signal is not always visible, and field equipment typically uses a lower digitization rate.

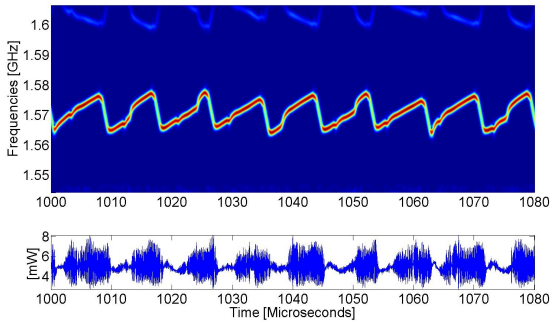

Figure 7 A jammer power spectra, with jumps in frequency. Continuous approximations to the above frequency behavior will likely be inaccurate and the phase estimate will likely be untrustworthy.

Because the phase estimated state θ is unlikely to be accurate, the phase state will be removed for the single jammer signal tracking case. The resulting measurements are further simplified by considering only the absolute value of the accumulations, as follows:

$$\begin{aligned} z(L) &= |I + (iQ)\,| \\ &= |\sum_{n=1}^{N} w(n)y(n)e^{-i\zeta(L,n,N)}| \\ &\quad \forall L \in [1, ..., N] \end{aligned} \tag{23}$$

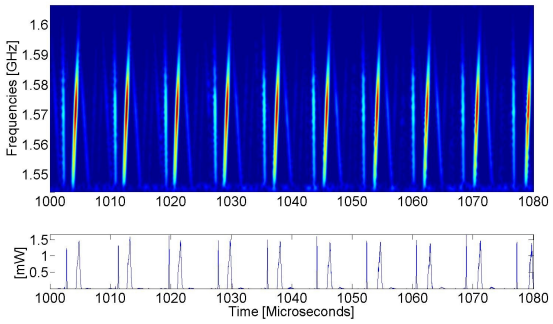

Figure 8 A jammer power spectra, the frequency span is much greater than that covered by the high--speed data sampling equipment (62.5 MHz complex samples). Phase estimates derived from data that spans multiple ramps will likely be erroneous.

The measurement model is defined similarly:

$$\begin{aligned} h(L) &= |I(\underline{x}) + (iQ(\underline{x}))\,| \\ &= |\sum_{n=1}^{N} w(n)\bar{y}(n, \underline{x}, f_{mix})e^{-i\zeta(L,n,N)}| \\ &\quad \forall L \in [1, ..., N] \end{aligned} \tag{24}$$

where $\bar{y}$ is the predicted RF sample obtained by propagating the state to the current measurement interval and evaluating Eq. 17 for the time at sample n. The I and Q dependencies on the non-state terms have been omitted for the sake of notational convenience. The Jacobian, or partial derivatives of the measurement model with respect to the model state, is required for Kalman filter estimation and is defined as:

$$\begin{aligned} H(L) &= \frac{\partial}{\partial \underline{x}}|I(\underline{x}) + (iQ(\underline{x}))\,| \\ &= \frac{\partial}{\partial \underline{x}}|\sum_{n=1}^{N} w(n)\bar{y}(n, \underline{x}, f_{mix})e^{-i\zeta(L,n,N)}| \\ &\quad \forall L \in [1, ..., N] \end{aligned} \tag{25}$$

The absolute value is defined as:

$$|a + (ib)| = \sqrt{(a^2) + (b^2)} \tag{26}$$

and the partial derivative as:

$$\frac{\partial}{\partial \underline{x}}|a + (ib)| = \frac{a\left(\frac{\partial a}{\partial \underline{x}}\right) + b\left(\frac{\partial b}{\partial \underline{x}}\right)}{|a + (ib)|} \tag{27}$$

In the case of the measurement model from Eq. 24:

$$\begin{aligned} a &= I\left(\underline{x}\right) \\ &\approx \frac{A}{2}\sum_{n=1}^{N} w(n) cos\left(\Theta(\underline{x}, t_n, f_{mix}) - \zeta\left(L, n, N\right)\right) \\ &= \frac{A}{2}\sum_{n=1}^{N} w(n) cos\left(\mu\left(L, n, N, f_{mix}, \underline{x}\right)\right) \\ &= \frac{A}{2}\sum_{n=1}^{N} w(n) cos\left(\mu\left(\underline{P}, \underline{x}\right)\right) \end{aligned} \tag{28}$$

$$\begin{aligned} b &= Q\left(\underline{x}\right) \\ &\approx \frac{A}{2}\sum_{n=1}^{N} w(n) sin\left(\Theta(\underline{x}, t_n, f_{mix}) - \zeta\left(L, n, N\right)\right) \\ &= \frac{A}{2}\sum_{n=1}^{N} w(n) sin\left(\mu\left(L, n, N, f_{mix}, \underline{x}\right)\right) \\ &= \frac{A}{2}\sum_{n=1}^{N} w(n) sin\left(\mu\left(\underline{P}, \underline{x}\right)\right) \end{aligned} \tag{29}$$

where $\underline{P}$ is a vector of parameters that has been introduced for notational convenience and includes L, n, N, andf_{mix}. The reason that Eq. 28 and 29 are only approximations is because the high frequency term, from the conversion of the products of the trigonometric terms to the sum of the terms with combined arguments, is assumed to sum to zero. The resulting partial derivatives of a and b are as follows:

$$\frac{\partial a}{\partial \underline{x}} = \begin{cases} \frac{a}{A}, \text{for } x = A \\ -\frac{A}{2}\sum_{n=1}^{N} w(n)\frac{\partial \Theta}{\partial \underline{x}} sin\left(\mu\left(\underline{P}, \underline{x}\right)\right), \text{otherwise} \end{cases} \tag{30}$$

$$\frac{\partial b}{\partial \underline{x}} = \begin{cases} \frac{b}{A}, \text{for } x = A \\ \frac{A}{2}\sum_{n=1}^{N} w(n)\frac{\partial \Theta}{\partial \underline{x}} cos\left(\mu\left(\underline{P}, \underline{x}\right)\right), \text{otherwise} \end{cases} \tag{31}$$

Eqs. 26–31 can be combined with the measurement model equations, Eq. 24 and 25, to complete the absolute value FFT measurement model.

Sometimes it might be preferable to use the power of an accumulation pair instead of just the absolute value, where the absolute value is effectively the square root of the power. In that case the measurements are defined using the same I and Q definitions from the FFT model:

$$\begin{aligned} z(L) &= \left[I + (iQ)\right]^* \left[I + (iQ)\right] \\ &= I^2 + Q^2 \\ &\forall L \in [1, ..., N] \end{aligned} \tag{32}$$

The measurement model is defined as follows:

$$\begin{aligned} h(L) &= I(\underline{x})^2 + Q(\underline{x})^2 \\ &\forall L \in [1, ..., N] \end{aligned} \tag{33}$$

with partial derivatives:

$$\begin{aligned} H(L) &= \frac{\partial}{\partial \underline{x}}\left[I(\underline{x})^2 + Q(\underline{x})^2\right] \\ &= 2\frac{\partial I(\underline{x})}{\partial \underline{x}} I(\underline{x}) + 2\frac{\partial Q(\underline{x})}{\partial \underline{x}} Q(\underline{x}) \\ &\forall L \in [1, ..., N] \end{aligned} \tag{34}$$

The partial derivatives of the accumulations with respect to the states are given by Eqs. 30 and 31.

Every RF sample includes some amount of sample noise, which eventually becomes accumulation noise. The noise statistics of the absolute value of the FFTs is complicated and only an approximation has been used in this paper. The results are similar to the noise statistics for the power measurements, which are as follows:

$$E[z(L)] \approx \left(I_{true}^2 + Q_{true}^2\right) + N\sigma^2 \tag{35}$$

$$\begin{aligned} cov(z(L)) &= E\left[\left(z(L) - E[z(L)]\right)\left(z(L) - E[z(L)]\right)^*\right] \\ &\approx 2N\sigma^2\left[I_{true}^2 + Q_{true}^2\right] + N^2\sigma^4 \end{aligned} \tag{36}$$

where I_{true} and Q_{true} are the accumulations if the signal contained zero noise.

One component of the measurement model that still needs to be addressed is the RF filter characteristics of the physical jammers and recording devices. The RF filters will attenuate different frequencies at different rates. The data recording equipment RF filter shape can be determined off-line. Off-line system identification is likely not possible for the jammers, because the equipment is not typically available before the first encounter in the field. It might still be possible to perform off-line system identification on several jammers and attempt to identify repeatable RF filter shapes among the surveyed jammers. Therefore, if the class of the jammer is known *a priori* then the RF filter shape can be assumed—it may provide sufficient accuracy for tracking purposes.

Alternatively, the RF filter shapes can be identified online, after signal detection and before the initialization of the tracking algorithms. There are many ways to do this, but one simple method of system identification would start by computing FFTs of the incoming signal at a rate that is faster than the chirp period. Then the maximum magnitude, or the average magnitude, of

each FFT bin would be computed over many chirp periods. The resulting RF filter's magnitude versus frequency bin output would contain the effects of both the jamming and recording equipment. That output could then be normalized. The result of the above procedure is an attenuation factor that can be included into each FFT computation in the measurement model.

It should be noted that the maximum time length of each accumulation is a function of the frequency rate of change of the jammer being tracked. A jammer that has a slower ramp rate will be able to use more samples before the frequency mismatch between the jammer's frequency and the FFT bin frequency causes the I-Q vector to complete a full rotation. A jammer with a fast rate of change will only be able to use a fewer number of samples.

If the user desires to track a very weak signal, or a signal that has a quick time-rate-of-change of frequency, then a different measurement model might be required. A candidate model that has similar characteristics of an FFT accumulation model is one that starts with FFTs, but then adds a linear frequency rate of change term. The ramping frequency term will allow for longer accumulations, and therefore a higher signal-to-noise ratio. However, there are several down-sides to this new proposed model. One down-side is that many ramp-rates would be needed until the actual ramp rate of the jammer being tracked is known with enough certainty that only one ramp rate could be used. A second down-side is that the measurement noise covariance matrix, R, may become more complicated. This type of accumulation may also face length limitations due to the NCO frequency passing outside of the Nyquist bounds of the system. Finally, the observability properties of this new model have not been investigated.

AD-HOC ACQUISITION OF A MODERATELY POWERFUL CHIRP-STYLE JAMMER SIGNAL

There are multiple ways to initialize a state for use in the developed chirp-style GPS jammer signal tracking filter. The method presented in this paper assumes a moderately strong jammer signal-to-noise ratio, i.e. the signal can be at least partially seen in FFTs computed at the receiver. The primary obstacle to a fast state initialization procedure in this system is the high dimensionality of the proposed state. If the 2-ramp polynomial parameterization model developed in this paper is used (without phase) then the state's dimension is fifteen. This is a state space that is likely too large to search in any brute-force manner in a small amount of time. Instead, the initialization procedure will attempt to extract some state values from autocorrelations and FFTs of the sampled data and then make some simplifying assumptions about the jammer to reduce the search dimensions to a more tractable size. The developed algorithm is intended to allow for the initialization of the state of a single jammer only.

The spectra of a GPS jammer from data recorded in a field campaign at White Sands Missile Range in 2012 is shown in Fig. 9. Several quantities are clearly dis-

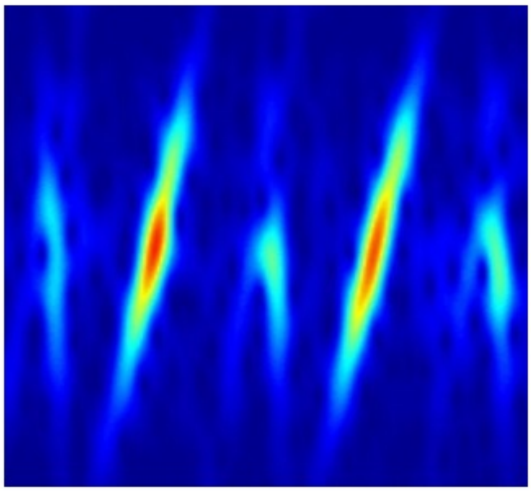

Figure 9 A jammer power spectra spanning 9 MHz on the y axis and approximately 20 μs on the x axis.

tinguishable from Fig. 9. These quantities include the period of the jammer, T, a time t_0 at which the signal passes through a given frequency f_0, and the linear coefficient of the ramp-up polynomial, c_1^u. The four quantities, T, f_0, t_0, and c_1^u are labeled in Fig. 10, and code can easily be written to automate the computation of those quantities. The period can also be computed more precisely by using autocorrelations of the raw data.

Fig. 10 seems to imply that the other states cannot reliably be initialized by directly inspecting the FFTs of two chirps, and an expanded initialization procedure should be developed.

The additional initialization procedure presented in this paper will collapse the remaining (large) search space by assuming a 1^{st} order linear polynomial chirp. This assumption is not strictly valid, but appears to be a reasonable low-order approximation of the many jammers in Ref. [8]. The benefit of the 1^{st} order linear polynomial chirp assumption is that it reduces the

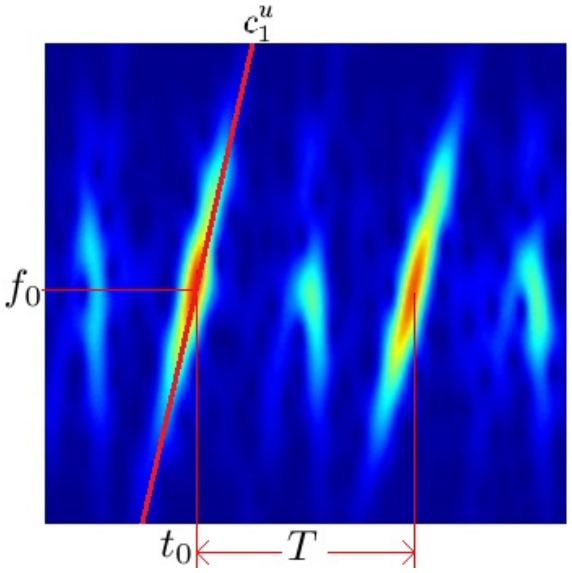

Figure 10 A jammer power spectra spanning 9 MHz on the y axis and approximately 20 μs on the x axis. Easily extractable quantities have been labeled.

number of polynomial coefficients that must be initialized to two. This lower-dimensional initialization can be considered a "rough" initialization.

It might appear that there are four states that must be determined to complete the rough initialization: f, c_1^d, t^u and t^d (the amplitude state can be scaled at the end of the initialization procedure). That is not the case. Coupling between states, the frequency f_0, and the time t_0, can be used to reduce the independent initialization dimension to two. The remaining two dimensions must be explored to determine the best rough state initialization. The first dimension is the span of the chirp. The span can be defined in terms of the frequency or time, because the two quantities are coupled through the c_1^u term. The frequency is selected in this work. The frequency span f_{span} is shown graphically for two distinct chirps in the Fig. 11. Where $f_{span,1}$ is the frequency span of the first chirp and $f_{span,2}$ is the frequency span of the second chirp.

The second dimension is the starting point of the chirp. The starting point can be defined in terms of either the initial frequency f or time of ramp up t^u, because the two quantities are coupled in a manner similar to the coupling in f_{span}. The frequency f is selected in this work, so that parallelism is preserved between the two search dimensions. The center frequency f_0 (or f_{L_1}), t_0, and c_1^u are provided earlier by direct computation on the FFT spectra, and they constrain the chirp to lie along a line in the frequency-time plot with slope c_1^u, and they prescribe an intersection point at (t_0,f_0). The quantities f_{span}, T, c_1^u, and the assumption of a 1st order linear polynomial chirp prescribe the ramp-down rate c_1^d. The starting frequency search is shown graphically in Fig. 12 for three different hypothesis starting chirp times and frequencies, f_1, f_2, and f_3, but identical states otherwise.

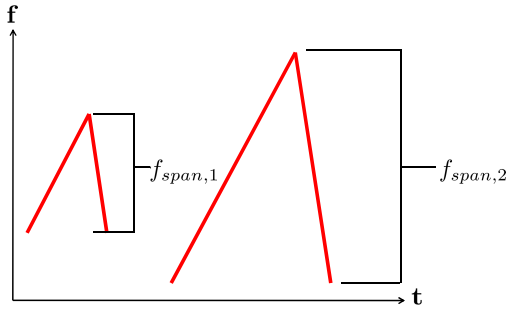

Figure 11 Frequency-time histories for two different values of the chirp frequency span, but otherwise using similar states. The chirps have been offset in time in the interest of clarity.

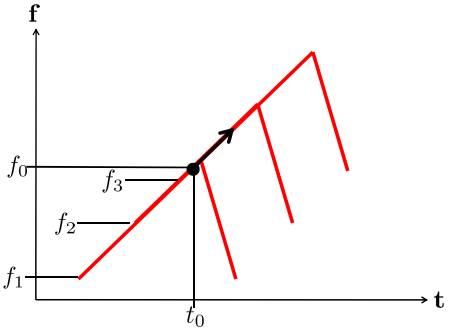

Figure 12 Frequency-time histories for three different values of the chirp frequency starting value, but similar states otherwise.

The many different state hypothesis $\underline{x}_h$that will result from searching the above two dimensions must be evaluated to determine how closely they fit the real data. There are many ways to define an appropriate ad-hoc "goodness of fit" metric. This paper's metric will be a cost function $J\left(\underline{x}_h, \underline{y}\right)$ that seeks to minimize the 2-norm of the frequency difference between the real data and the hypothesis data. The actual metric considers

the FFT frequency bin spacing distance between peaks of FFTs computed on real data and those computed on the hypothesis data. The formal definition of the scalar cost function is as follows:

$$J\left(\underline{x}_h, \underline{y}\right) = ||\underline{d}\left(\underline{x}_h, \underline{y}\right)||_2 \tag{37}$$

where $\underline{d}$ is a vector of FFT bin distances between the maximum powered accumulation in the FFTs of the real data and the hypothesis data.

Fig. 13 shows a possible entry of d, where the vector $\underline{d}$ is composed of many such d values, and corresponding FFTs on sequential data. In Fig. 13 d is the single distance between the peaks of the real data, at bin number 95, and the hypothesis data, at bin number 175, and is approximately 80 bins. Many computations of the form shown in Fig. 13 are combined to cover the full chirp period of the jammer as shown in Fig. 14. In fact, the use of several chirps seems to improve the 1st order linear frequency polynomial fit results by averaging to reduce the effect of measurement noise.

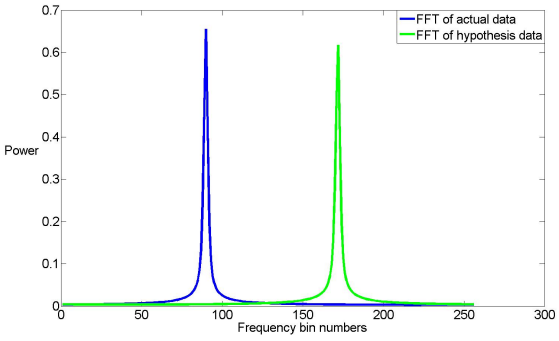

Figure 13 An example plot of FFT power versus FFT bin number for real data (blue) and for hypothesis data (green).

The final step in the initialization procedure is the "fine" initialization. The fine initialization procedure expands the rough initialization state corresponding to the best 1st order linear polynomial chirp approximation to a full polynomial. The above rough initialization procedure should produce a state estimate that is within the pull-in range of a standard nonlinear Maximum Likelihood Estimator (MLE) that allows full polynomial variation on the frequency ramps. The Square Root Information (SRI) formulation of the MLE cost function is shown below:

$$J\left(\underline{x}, \underline{y}\right) = \left[R_a^{-T}\left(\underline{z}\left(\underline{y}\right) - h\left(\underline{x}\right)\right)\right]^T \left[R_a^{-T}\left(\underline{z}\left(\underline{y}\right) - h\left(\underline{x}\right)\right)\right] \tag{38}$$

where $\underline{z}$ is the vector of FFT measurements computed from the data $\underline{y}$, h is the FFT measurement model of

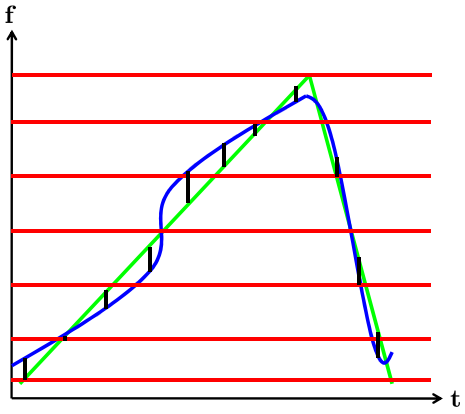

Figure 14 Graphical representation of a 1st order linear polynomial frequency approximation to a higher order chirp. FFTs of the real data are in blue, hypothesis data in green, the frequency bin distances between the two data (before rounding to the nearest frequency bin) are in black, and the the FFT frequencies used to calculated the entries of $\underline{d}$ are in red.

Eq. 24, and R_a is the Cholesky factorization, or matrix square root, of the measurement noise covariance matrix R. Where R is defined as:

$$R = R_a^T * R_a = E\left[\left(\underline{\nu} - E\left[\underline{\nu}\right]\right)\left(\underline{\nu} - E\left[\underline{\nu}\right]\right)^T\right] \tag{39}$$

where $\underline{\nu}$ is the measurement noise vector. The measurement noise statistics are similar to those in the previous section on jammer observables.

SINGLE-JAMMER KALMAN FILTER SIGNAL TRACKING

The state, dynamics model, measurement model, and initialization procedure can be combined to produce a Kalman filter to track the RF signal of a chirp-style GPS jammer. The data used in this section comes from two different sources. The first source is data originally collected for use in Ref. [8], where many GPS jammer's signals were recorded using complex sampling at a rate of 62.5 MHz. The second source is the data originally collected for use in Ref. [15], where many GPS jammer's signals were recorded at approximately 8 or 9 MHz (depending on the file) during a testing event at WSMR sponsored by the DHS.

The first set of Kalman filter tracking results considers three different jammers in a laboratory environment. The jammers progress from a nearly perfectly linear 1st order polynomial jammer behavior to the jammer

that was least like a 1st order jammer out of all the jammer signatures collected for Ref. [8]. Each laboratory jammer is initialized *not* using the previously developed acquisition procedure. The initialization was performed by manually inspecting an FFT surface plot of the jammer data and then guessing a linear 1st order polynomial state. This particular initialization was selected so that the initial state provided to the filter had some significant error, as was likely not to be the case with the previously developed acquisition procedure. The error was introduced to demonstrate some of the filter convergence properties. A full analysis of the filter convergence range is beyond the scope of this paper. This paper's ad-hoc acquisition procedure has been used on the WSMR data shown later.

The results of the most benign jammer, or that which has a frequency behavior that most closely matches the linear 1st order frequency polynomial model, is shown in Fig. 15. The Kalman filter's frequency estimate at each sample time is plotted on top of a surface plot comprised of FFTs of the real data. Despite the initial error, the filter quickly determines the correct state of the jammer, and, although it is not shown, the filter state estimates completely settle within 10 chirp periods.

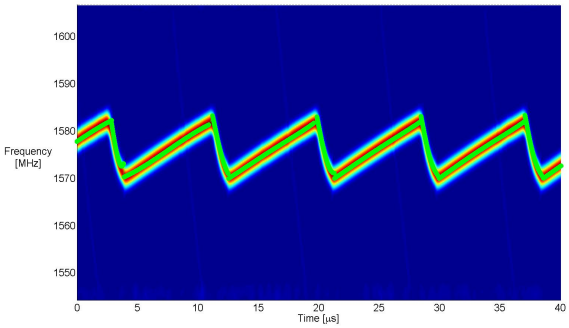

Figure 15 Spectra of the first GPS jammer in a laboratory environment, with the Kalman filter frequency state estimates overlaid in green.

The results of the second jammer are shown in Fig. 16. The second jammer does not have as close of a match to the linear 1st order frequency polynomial behavior as the previous jammer, as can be seen in the curve of the frequency up-ramp. The initial error can be seen in the jump of the frequency estimate at the bottom of the first chirp. The Kalman filter determines the state of the jammer in less than two chirp periods.

The results of the least benign jammer, or that which has a frequency behavior that least closely matches the linear 1st order frequency polynomial model, is shown in Fig. 17. Despite the fact that the jammer has a different ramp behavior for each chirp the Kalman filter is able to modify its state estimate every ramp to closely follow the jammer's actual frequency.

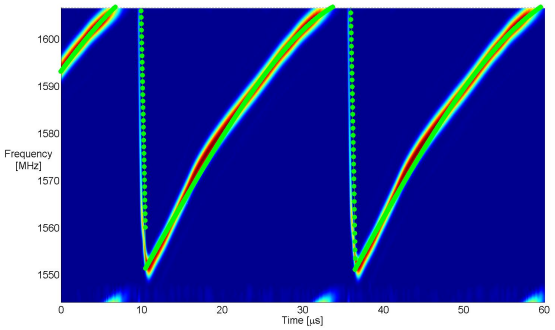

Figure 16 Spectra of the second GPS jammer in a laboratory environment, with the Kalman filter frequency state estimates overlaid in green.

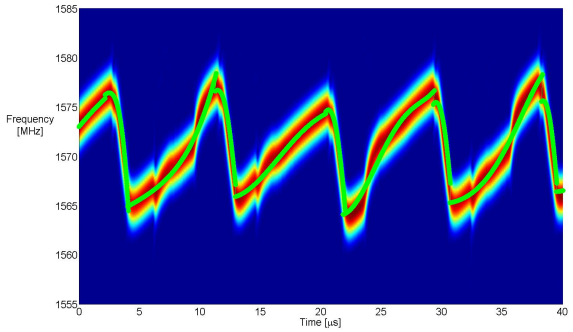

Figure 17 Spectra of the third GPS jammer in a laboratory environment, with the Kalman filter frequency state estimates overlaid in green.

The above results will naturally degrade with increased noise. The motivation for testing the filter with more noise is to understand how the jammers can be tracked in a real-world scenario. Instead of testing with more noise the filter will be analyzed using actual data from a real-world scenario at WSMR.

The second set of Kalman filter tracking results considers one jammer at two different distances from the receiver station. The first set of results is shown in Fig. 18, when the jammer is approximately 50 meters from the recording station. The state is initialized using the automated acquisition procedure mentioned previously and the filter is able to track the signal very accurately from the beginning of the data set.

In Fig. 19, the Kalman filter is applied to the same jammer as in Fig. 18, but now the jammer is approx-

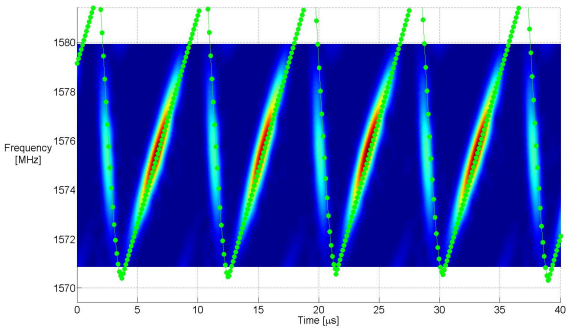

Figure 18 Spectra of a GPS jammer at WSMR when it is very close, with the Kalman filter frequency state estimates overlaid in green.

imately 1.8 kilometers away from the receiver station. Despite the fact that the jamming signal is only barely visible above the noise the acquisition procedure is able to initialize the state and the filter is able to track the signal. It should be noted that the automated acquisition procedure required several initialization attempts, discarding several μs of data after each failed initialization, to produce a good state estimate for signal tracking. Methods of determining when the signal has failed initialization are beyond the scope of this paper, but a common test is to consider the measurement residuals after the initialization procedure concludes.

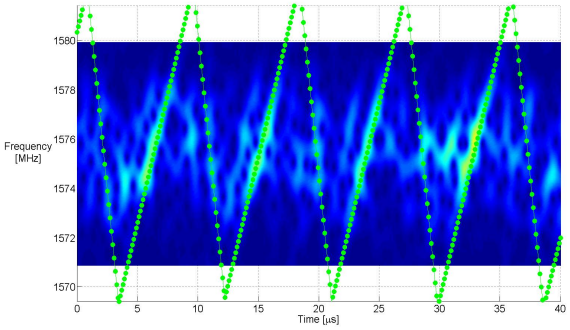

Figure 19 Spectra of a GPS jammer at WSMR when it is far away, with the Kalman filter frequency state estimates overlaid in green.

MULTI-JAMMER KALMAN FILTER SIGNAL TRACKING

The jammer models must first be extend to handle multiple signals before the signal tracking problem can be addressed. The dynamics of each jammer are independent of other jammers and result in the following uncoupled dynamics equation:

$$\underline{\underline{x}}_{k+1} = \underline{\underline{\Phi}}\ \underline{\underline{x}}_k + \underline{\underline{\Gamma}}\ \underline{\underline{v}}_k \tag{40}$$

where the entries of Eq. 40 are defined as follows for the case of two jammers:

$$\begin{aligned}
\underline{\underline{x}} &= \begin{bmatrix} \underline{x}_1 \\ \underline{x}_2 \end{bmatrix} \\
\underline{\underline{v}} &= \begin{bmatrix} \underline{v}_1 \\ \underline{v}_2 \end{bmatrix} \\
\underline{\underline{\Phi}} &= \begin{bmatrix} \Phi_1 & 0 \\ 0 & \Phi_2 \end{bmatrix} \\
\underline{\underline{\Gamma}} &= \begin{bmatrix} \Gamma_1 & 0 \\ 0 & \Gamma_2 \end{bmatrix}
\end{aligned}$$

where $\underline{x}_1$ is the state of the first jammer, $\underline{x}_2$ is the state of the second jammer, and the other terms are defined in a similar manner in conjunction with Eq. 7–13. Eq. 40 can be extended to handle an arbitrary number of jammers by appending the new jammer's state to the current state vector, appending the new process noise vector to the current process noise vector, and then modifying the matrices $\underline{\underline{\Phi}}$ and $\underline{\underline{\Gamma}}$ appropriately.

The measurement model is slightly more complicated. The signal received at the recording station is a linear combination of the two, or more, jammer signals. The combined signal model is as follows:

$$\begin{aligned}
\bar{y}\left(t_i, \underline{\underline{x}}, f_{mix}\right) &= \bar{y}_1\left(t_i, \underline{x}_1, f_{mix}\right) + \bar{y}_2\left(t_i, \underline{x}_2, f_{mix}\right) \\
&= \frac{A_{1,i}}{2} * cos\left(\Theta\left(\underline{x}_1, t_i, f_{mix}\right)\right) \\
&+ \frac{A_{2,i}}{2} * cos\left(\Theta\left(\underline{x}_2, t_i, f_{mix}\right)\right)
\end{aligned} \tag{41}$$

where $\bar{y}_1$ and $\bar{y}_2$ are defined by Eq. 17. Eq. 41 can be extended to an arbitrary number of jammers by adding the new signals in the same manner that $\bar{y}_2$ was added to $\bar{y}_1$.

The resulting FFT measurement model has the same fundamental form as before:

$$\begin{aligned}
h(L) = |\sum_{n=1}^{N} w(n)\bar{y}(n, \underline{\underline{x}}, f_{mix})e^{-i\zeta(L,n,N)}| \\
\forall L \in [1, ..., N]
\end{aligned} \tag{42}$$

FFT operators have the convenient property that they are linear and can therefore be distributed among the multiple incoming signals. The resulting measurement

model equation has the following form:

$$h(L) = |\sum_{n=1}^{N} w(n)\bar{y}_1(n, \underline{x}_1, f_{mix})e^{-i\zeta(L,n,N)} + ... \sum_{n=1}^{N} w(n)\bar{y}_2(n, \underline{x}_2, f_{mix})e^{-i\zeta(L,n,N)}| \quad (43)$$
$$\forall L \in [1, ..., N]$$

The partial derivatives behave similarly and have been omitted for the sake of brevity.

Multiple incoming jamming signals experience an additional effect that is not normally seen in single jamming signals: interference. The two signals are frequency modulated sinewaves that will likely have very similar frequencies at some point during signal tracking. When two sine waves with similar frequencies are combined they begin to constructively and destructively interfere with each other. The factor that determines whether the interference is constructive or destructive is the relative phase of the two signals. Therefore, the signal phase that was neglected earlier in this paper's filter development becomes vitally important. The measurement model could be modified again to allow the phase of both signals to be states, but it is difficult to track both signal's phases accurately enough to predict the type and level of interference. Additionally, it is the relative phase between the jammers that determines the type and extent of the interference. Therefore, at each time step the relative phase θ_{rel} of the jammers is added as a parameter and the measurements are optimized over it. The optimization cost function is the same as the earlier maximum likelihood cost function, but for a single measurement set:

$$J\left(\theta_{rel}, \underline{x}, \underline{y}\right) = \left[R_a^{-T}\left(\underline{z}\left(\underline{y}\right) - h\left(\underline{x}, \theta_{rel}\right)\right)\right]^T \left[R_a^{-T}\left(\underline{z}\left(\underline{y}\right) - h\left(\underline{x}, \theta_{rel}\right)\right)\right] \quad (44)$$

Determining the optimal relative phase θ_{rel} from Eq. 44 is easily accomplished by evaluating multiple phases on the unit circle and selecting the one with the lowest cost. It was found that a grid of 40 points spaced evenly on the unit circle determined a relative phase value to an appropriate resolution for tracking two jamming signals.

The above algorithm was tested on laboratory data. A multi-jammer data set was developed by adding together the samples from two different jammer data sets. The states of each jammer were initialized separately on the individual data sets before combination because an efficient multi-jammer initialization algorithm is beyond the scope of this current work. The spectra of the combined jammer samples is shown in Fig. 20. The Kalman filter frequency estimates of the two jammers are overlaid in green and red in Fig. 21.

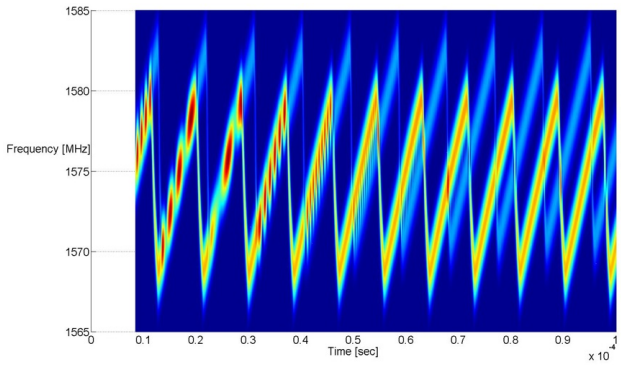

Figure 20 Spectra of two GPS jammers in a laboratory environment.

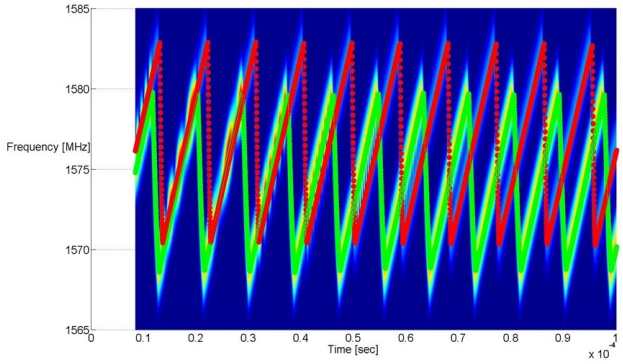

Figure 21 Spectra of two GPS jammers in a laboratory environment, with the Kalman filter frequency estimates overlaid in red and green.

The multi-jammer signal tracking work has been shown to work on real jamming data in a laboratory setting. A real-world test has not yet been completed, and it is not known how many jammers can reasonably be tracked using the developed algorithms. It should be noted that the optimization in Eq. 44 increases in dimensionality by one every time another jammer is added to the incoming signal, effectively slowing down the algorithm more than might normally be expected from the simple state dimension increase. It should also be noted that significant tuning of the measurement and process noise was required to allow the filter to track the combined two-jammer signal. The results of the multi-signal tracking algorithm would likely improve, and be more stable to noise parameters, if a filter architecture that is different from the Extended Kalman filter was implemented. A filter architecture that might work more reliably is the particle filter or the Gaussian mixture filter.

SUMMARY AND CONCLUSIONS

This paper investigated one method of acquiring and one method of tracking various chirp-style GPS jammers. A signal model that uses polynomial descriptions of frequency versus time was developed and its utility for jammer signal tracking and classification was discussed. Two slightly different observation models were derived. An ad-hoc chirp-style signal acquisition method that used FFTs of a moderately strong jamming signal was also developed. The various algorithm components were combined to produce a signal tracking Kalman filter. The developed Kalman filter was verified by application on three sets of laboratory data and on two sets of field data. One GPS jammer was even acquired and tracked at a distance of approximately 1.8 kilometers. The developed models and algorithms were expanded to consider multiple jammers simultaneously, and the algorithms were tested on one set of multi-jammer laboratory data.

ACKNOWLEDGEMENTS

The authors would like to express their thanks to the Department of Homeland Security for their arrangement and sponsorship of the jamming event at White Sands Missile Range and to the 746th Test Squadron for their part in the jamming event.

The authors would also like to express their thanks to the following members of the UT/Austin Radionavigation Laboratory for helping to staff the receiver stations at White Sands Missile Range: Todd Humphreys, Daniel Shepard, and Reese Shetrone.

REFERENCES

[1] Arthur, C., "Car thieves using GPS 'jammers'," Monday 22 February 2010, http://www.guardian.co.uk/technology/2010/feb/22/car-thieves-using-gps-jammers.

[2] Anon., "National PNT Advisory Board comments on Jamming the Global Positioning System - A National Security Threat: Recent Events and Potential Cures," November 2010.

[3] Misra, P. and Enge, P., *Global Positioning System: Signals, Measurements, and Performance*, Ganga-Jamuna Press, Lincoln, Massachusetts, 2nd ed., 2006, pp. 53–58.

[4] Spilker, J. and Natali, F., *Global Positioning System: Theory and Applications Volume 1*, American Institute of Aeronautics and Astronautics, Inc., Washington, DC, 1st ed., 1996.

[5] Hyde, J., "GPS Tracking Device Leads to Arrest of Evanston Bank'," August 18, 2011, http://www.fbi.gov/chicago/press-releases/2011/gps-tracking-device-leads-to-arrest-of-evanston-bank-robbery-suspect.

[6] Pullen, S. and Gao, G., "GNSS Jamming in the Name of Privacy," *Inside GNSS*, Vol. 7, Mar./Apr. 2012.

[7] Mitch, R. H., Dougherty, R. C., Psiaki, M. L., Powell, S. P., O'Hanlon, B. W., Bhatti, J. A., and Humphreys, T. E., "Innovation Column: Know Your Enemy," *GPS World*, January 2012.

[8] Mitch, R. H., Dougherty, R. C., Psiaki, M. L., Powell, S. P., O'Hanlon, B. W., Bhatti, J. A., and Humphreys, T. E., "Signal Characteristics of Civil GPS Jammers," *Proceedings of the ION GNSS 2011*, Sept. 20-23, 2011, pp. 1907–1919, Portland, OR.

[9] Kraus, T., Bauernfeind, R., and Eissfeller, B., "Survey of In-Car Jammers - Analysis and Modeling of the RF signals and IF samples (suitable for active signal cancellation)," *Proceedings of the ION GNSS 2011*, Sept. 20-23, 2011, pp. 430–435, Portland, OR.

[10] Aloi, D. and Steffes, A., "Vehicle Impact on Personal Privacy Device (PPD) Performance," *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012)*, Sept. 2012, pp. 3558–3562, Nashville, TN.

[11] Gao, G., Gunning, K., Walter, T., and Enge, P., "Impact of Personal Privacy Device for WAAS Aviation Users," *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012)*, Sept. 2012, pp. 235–241, Nashville, TN.

[12] Potter, B., Shallberg, K., and Grabowski, J., "Personal Privacy Device Interference in the WAAS," *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012)*, Sept. 2012, pp. 2868–2874, Nashville, TN.

[13] Hegarty, C., "Considerations for GPS Spectrum Interference Standards," *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012)*, Sept. 2012, pp. 2921–2929, Nashville, TN.

[14] Bhatti, J., Humphreys, T. E., and Ledvina, B. M., "Development and Demonstration of a TDOA-

Based GNSS Interference Signal Localization System," *Proceedings of the IEEE/ION PLANS Conference*, Myrtle Beach, SC, April 2012.

[15] Mitch, R. H., Psiaki, M. L., O'Hanlon, B. W., Powell, S. P., and Bhatti, J. A., "Civilian GPS Jammer Signal Tracking and Geolocation," *Proceedings of the ION GNSS 2012*, Sept. 18-21, 2012, pp. 2901–2920, Nashville, TN.

[16] Fontanella, D., Bauernfeind, R., and Eissfeller, B., "In-Car GNSS Jammer Localization with a Vehicular Ad-Hoc Network," *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012)*, Sept. 2012, pp. 2885–2893, Nashville, TN.

[17] Trinkle, M., Cetin, E., Thompson, R., and Dempster, A., "Interference Localisation within the GNSS Environmental Monitoring System (GEMS) - Initial Field Test Results," *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012)*, Sept. 2012, pp. 2930–2939, Nashville, TN.

[18] Xu, Z. and Trinkle, M., "Weak GPS Interference Direction of Arrival Estimation Using GPS Signal Cancellation," *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012)*, Sept. 2012, pp. 2940–2945, Nashville, TN.