

Detection, Identification, Location, and Remote Sensing Using SAW RFID Sensor Tags

Richard J. Barton, Timothy F. Kennedy, Robert M. Williams, Patrick W. Fink, and Phong H. Ngo

NASA Johnson Space Center
Houston, TX

1. Introduction

The Electromagnetic Systems Branch (EV4) of the Avionic Systems Division at NASA Johnson Space Center in Houston, TX is studying the utility of surface acoustic wave (SAW) radio-frequency identification (RFID) tags for multiple wireless applications including detection, identification, tracking, and remote sensing of objects on the lunar surface, monitoring of environmental test facilities, structural shape and health monitoring, and nondestructive test and evaluation of assets. For all of these applications, it is anticipated that the system utilized to interrogate the SAW RFID tags may need to operate at fairly long range and in the presence of considerable multipath and multiple-access interference. Towards that end, EV4 is developing a prototype SAW RFID wireless interrogation system for use in such environments called the Passive Adaptive RFID Sensor Equipment (PARSEQ) system. The system utilizes a digitally beam-formed planar receiving antenna array to extend range and provide direction-of-arrival information coupled with an approximate maximum-likelihood signal processing algorithm to provide near-optimal estimation of both range and temperature. The system is capable of forming a large number of beams within the field of view and resolving the information from several tags within each beam. The combination of both spatial and waveform discrimination provides the capability to track and monitor telemetry from a large number of objects appearing simultaneously within the field of view of the receiving array.

In this paper, we will consider the application of the PARSEQ system to the problem of simultaneous detection, identification, localization, and temperature estimation for multiple objects. We will summarize the overall design of the PARSEQ system and present a detailed description of the design and performance of the signal detection and estimation algorithms incorporated in the system. The system is currently configured only to measure temperature (jointly with range and tag ID), but future versions will be revised to measure parameters other than temperature as SAW tags capable of interfacing with external sensors become available. It is anticipated that the estimation of arbitrary parameters measured using SAW-based sensors will be based on techniques very similar to the joint range and temperature estimation techniques described in this paper.

2. SAW RFID Background

SAW RFID technology differs from more conventional integrated-circuit (IC) based RFID technology in many significant ways. For example, SAW RFID tags do not need to rectify incident electromagnetic power arriving at the tag. Instead, they modulate and re-radiate the interrogation signal directly using a series of reflectors printed on the surface of the device. The design and operation of a typical SAW RFID tag is illustrated in Figure 1.

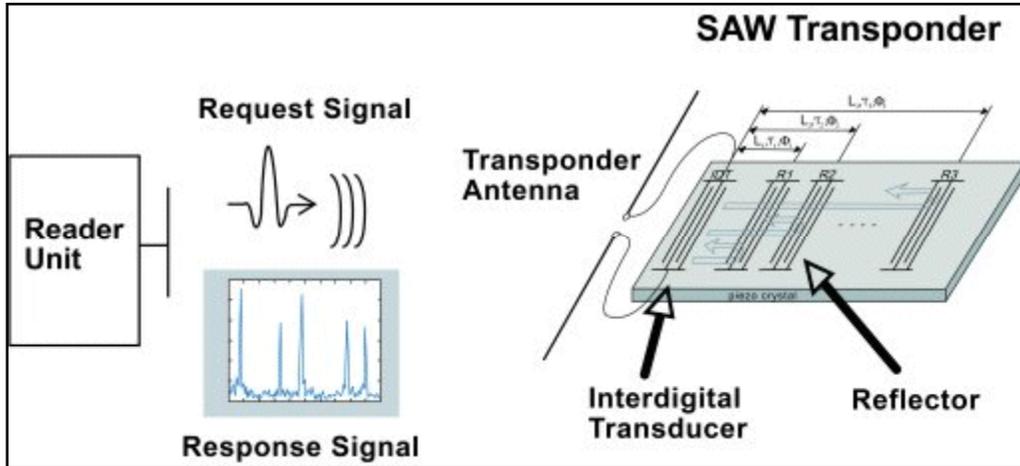


Figure 1. Design and operation of a typical SAW RFID tag.

The principle benefit of this approach is that SAW RFID tags are capable of operating at much lower interrogator transmit power, and/or are capable of much greater range for a given transmit power than IC-based tags [1]. In addition, SAW RFID tags have the capability to incorporate certain types of sensor telemetry along with the identification information. For example, the impulse response of a SAW RFID tag is highly sensitive to temperature. The temperature of the tag can be estimated directly by using a correlation operation to measure the time dilation (or contraction) of the impulse response. Finally, SAW RFID tags are very robust and can tolerate extreme conditions that would render IC-based tags inoperable. For example, they can withstand temperatures from cryogenic to over 310 C, they are inherently immune to ionizing radiation, and they have survived shock levels exceeding 1000 g.

The temperature of a SAW RFID tag can be estimated by direct measurement of the time dilation (or contraction) of the tag impulse response. In particular, measurement of the time dilation of the impulse response at an arbitrary temperature relative to the response at a known reference temperature (usually 0° C) constitutes an observation of the *temperature coefficient of delay* (TCD) for the tag at its current temperature. Here, the term TCD refers to the mathematical function of temperature that quantifies the relationship between the relative time dilation of the tag response and the temperature of the tag, with respect to a fixed reference temperature. Although the TCD can theoretically be determined from the piezo-electric properties of the crystalline material used to manufacture the tag, it is more common (and probably more accurate) to estimate it experimentally.

As an example, consider the waveform presented in Figure 2 below. This waveform represents the magnitude of a normalized version of the impulse response from a SAW RFID tag at a temperature of -73° C. The tag illustrated here is an example of the 40-bit Global SAW tag manufactured by RF SAW, Inc. [2]. There are a total of 22 pulses of slightly varying amplitude spaced intermittently over approximately a 1.6 μ s time period after an initial delay of approximately 1.2 μ s. The exact spacing of the pulses, excluding the first and last pulse (which are used for synchronization) along with the phase of each pulse (which is not visible in Figure 2) encodes the 40-bit ID of the individual tag.

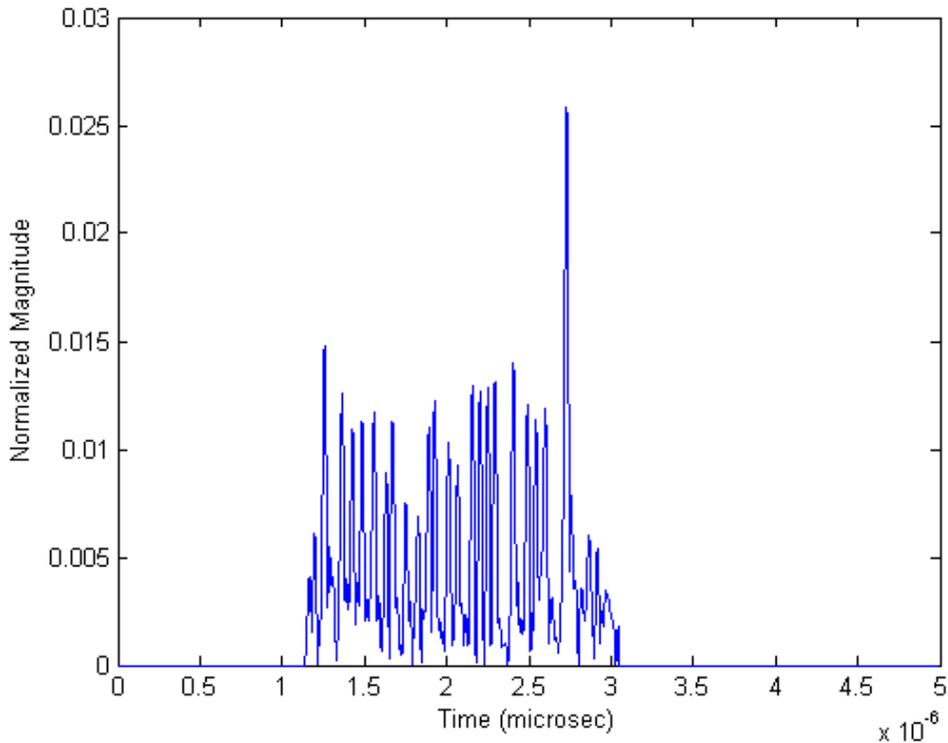


Figure 2. Magnitude of normalized tag impulse response at -73°C .

To illustrate the behavior of the RFID waveform a little more clearly, consider the waveforms illustrated in Figures 3 through 5. Figure 3 shows a detailed representation of one of the pulses from Figure 2 plotted showing both amplitude and phase information. As this figure indicates, at -73°C , each pulse lasts approximately 45 ns and oscillates at a frequency of approximately 2.44 GHz. Similarly, Figure 4 shows a detailed representation of two pulses from the waveform depicted in Figure 2 shifted to baseband with a one-sided bandwidth of 80 MHz. In this case, the envelope of the signal is superimposed over the true signal waveform. This baseband equivalent waveform captures all of the information necessary to estimate time delay precisely from the pulses and eliminates the high-frequency oscillatory behavior that obscures some of the detail in the signal. Once again, the duration of each pulse is approximately 45 ns but the frequency of oscillation has been reduced to approximately 40 MHz.

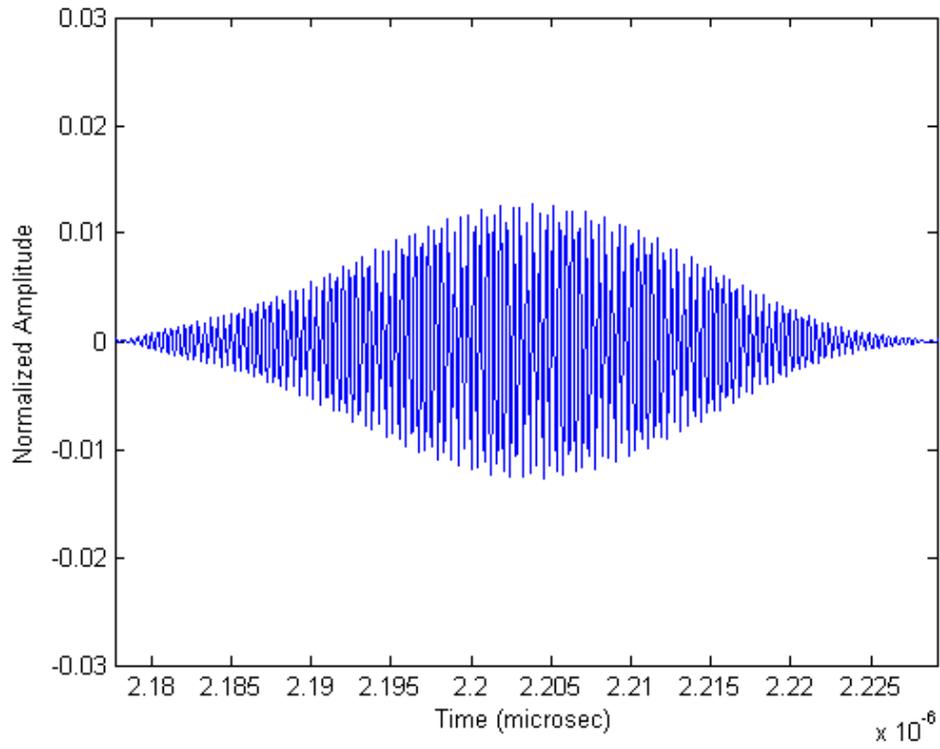


Figure 3. Detail of pulse from Figure 2.

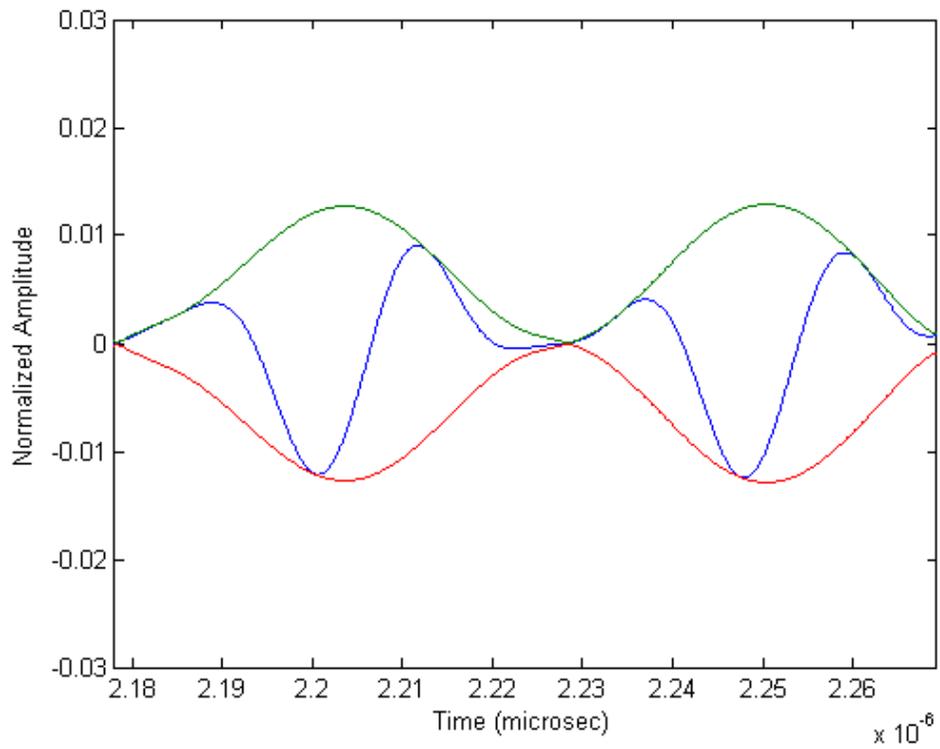


Figure 4. Detail of two baseband pulses at -73°C .

Finally, Figure 5 shows the magnitude of all 22 pulses in the RFID response waveform at both -73° C (blue) and 21° C (green). It is clear from the figure that the higher temperature waveform has been dilated with respect to the lower temperature waveform. In fact, the dilation factor is approximately 1.0068, which corresponds to a change in dilation factor of approximately 7.2×10^{-5} per degree C. As it turns out, the TCD for these tags is very close to linear over large temperature ranges and 7.2×10^{-5} represents a very good approximation to the slope of the corresponding line in the region from at least -100° C to 100° C.

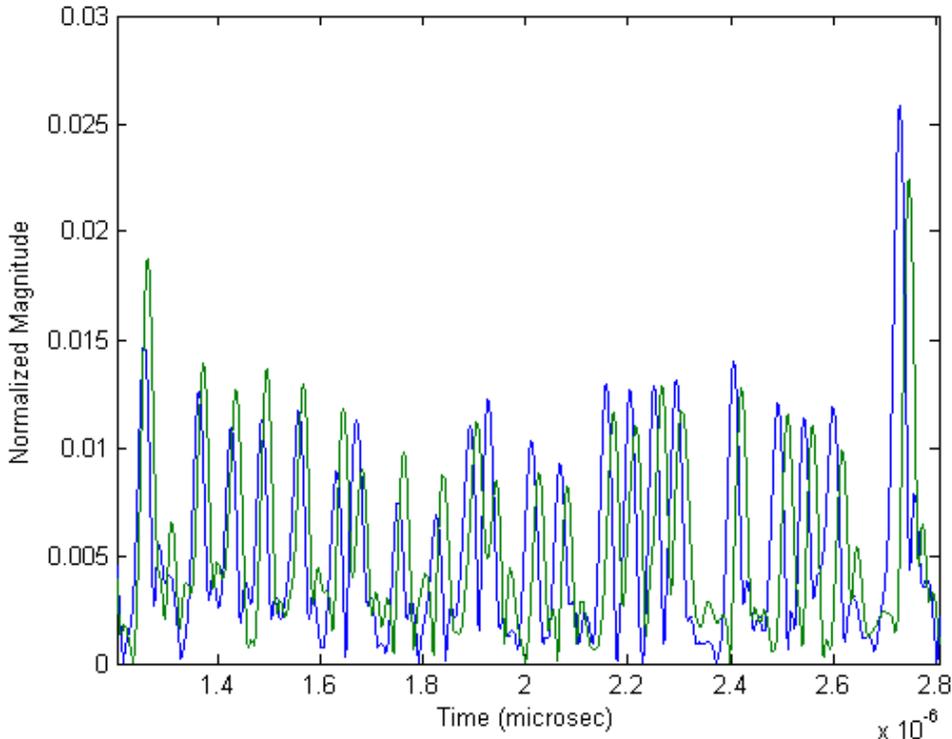


Figure 5. All 22 pulses of Global SAW tag impulse response magnitude at -73° C (blue) and 21° C (green).

3. PARSEQ System Description

The hardware architecture for the PARSEQ system is illustrated in Figure 6. The components of the system are as follows.

- *National Instruments (NI) Chassis and Peripherals* – The NI chassis contains a CPU board with USB and 100-Mbs Ethernet ports and a GPIB network interface board. The CPU board is used to control the operation of the overall system, perform all signal processing and beamforming operations, and provide an operator interface via an attached keyboard, mouse, and monitor. The GPIB network interface board supports communication between the CPU board and an RF switch matrix.
- *Agilent Network Analyzer (ENA)* – The ENA is used to measure the frequency response of each of the 72 transmit-receive channels supported by the system. The ENA is operated in stepped-CW mode and measures the S12 channel response; that is, to

measure the frequency response at any particular frequency, a sinusoidal signal is transmitted continuously from Channel 1 of the ENA over a transmit antenna array while the amplitude and phase of the received sinusoidal response is measured on Channel 2 of the ENA through a single antenna element on a receive antenna array routed through the RF switch matrix. The ENA is connected to the NI Chassis via the Ethernet LAN, which carries both control and data traffic between the CPU board in the chassis and the ENA. The measurement of channel frequency response using the ENA is described in more detail in Appendix A of this paper.

- *Transmit Antenna Array* – The transmit antenna array consists of 8 discrete circular patch antenna elements connected to the input via a 8-way power divider. The transmit antenna array is connected to channel 1 of the ENA via a low-loss, phase-coherent coaxial interconnect cable.
- *Receive Antenna Array* – The receive antenna array consists of 72 discrete circular patch antenna elements. Each antenna element is individually connected to one of the input channels on the RF switch matrix via a low-loss, phase-coherent coaxial interconnect cable.
- *Keithley RF Switch Matrix* – The RF switch matrix is used to connect individual elements of the receive antenna array to channel 2 of the ENA via low-loss, phase-coherent interconnect cables. The switch matrix is controlled by the CPU board in the NI chassis over the GPIB bus.

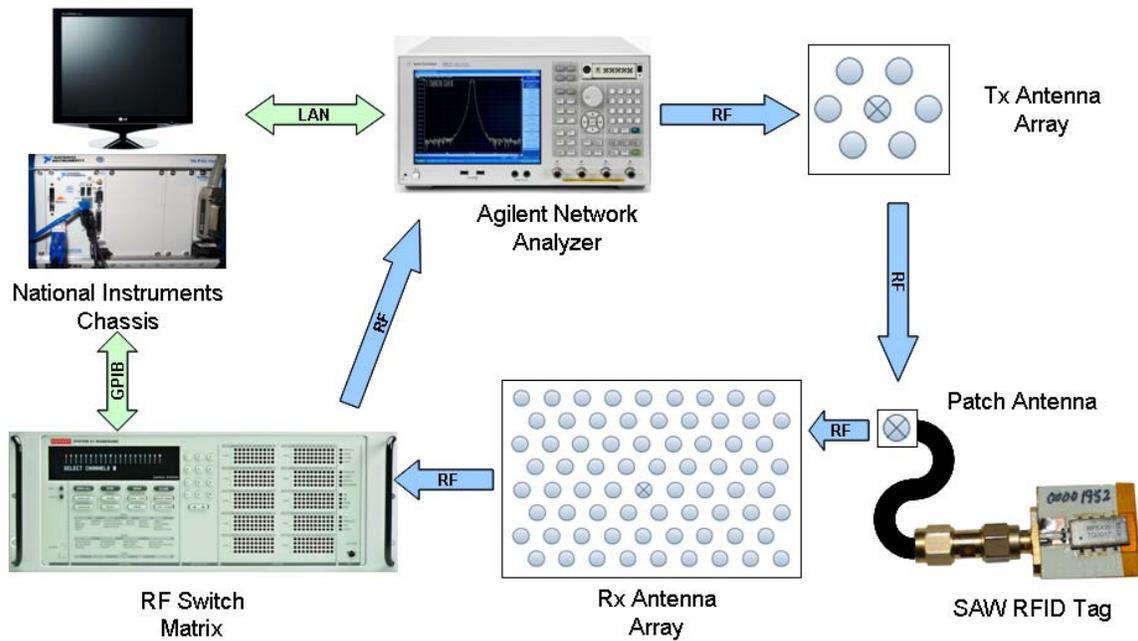


Figure 6. Hardware architecture for EV4 PARSEQ SAW RFID interrogator system.

The software architecture of the system, as currently configured, is illustrated in Figure 7. All software components of the system are hosted on the NI chassis CPU board. The software components of the system are as follows.

- *LabView Script* – The entire system is controlled by the LabView script. To acquire a single scan of data, the script sequentially sets the RF switch matrix to connect a single data stream from the receive array (i.e., a single antenna element) to Channel 2 of the ENA, commands the ENA to collect a sweep of the sampled frequency response for that channel, writes the frequency sweep data to a disk file, and then initiates the next step in the sequence. After collecting one full scan of data, the script passes control to an adaptive beamforming algorithm, which computes the desired beam coefficients from the sampled frequency-domain data and/or forms the desired beams using the computed beam coefficients. After forming all desired beams of frequency-domain data, the beamformer writes the beam data to a disk file and passes control back to the LabView script, which then passes control to a detection and estimation algorithm. The detection and estimation algorithm processes the beam data sequentially to estimate the range and temperature for each SAW RFID sensor tag detected in each individual beam. The range estimates along with the pre-determined azimuth and elevation associated with each beam are used to form a three-dimensional location estimate for each detected tag. After processing all beams, the detection and estimation algorithm writes the location and temperature estimates for all tags to a disk file and passes control back to the script, which displays the location/temperature estimates then continues the cycle by collecting a new scan of data.
- *Beamforming Algorithm* – The beamformer operates in both an adaptive mode and a feed-forward mode. In the adaptive mode, the algorithm reads in the frequency-domain data for all of the receive array channels, computes and displays the spatial power distribution across the receive array, computes and stores the beam coefficients, azimuth, and elevation corresponding to a pre-selected number of the strongest beams, and displays the time-domain response for each selected beam. In feed-forward mode, the beamformer uses the previously computed and stored beam coefficients to form the desired beams in the frequency domain and writes the beam data to the disk file for use by the detection and estimation algorithm. In the current implementation, the beamformer also performs some preliminary signal conditioning operations prior to computing and storing the beam data. In future implementations, the signal conditioning operations will be moved to a separate signal conditioning block in order to facilitate performance tuning of the system.
- *Detection and Estimation Algorithm* – The detection and estimation algorithm processes the beam-formed data and returns estimates of location and temperature for all tags detected in each of the beams. All processing occurs in the frequency domain, and tags are detected by correlating the received signal in each beam with a reference data base of tag *templates* measured at a pre-determined reference temperature. Only tags represented in the template data base can be detected by the current version of the software. It is hoped that future versions of the software will be modified to operate based on modeled rather than measured tag templates and to enable detection of all tags within the constellation of possible templates rather than a small subset of pre-determined templates.

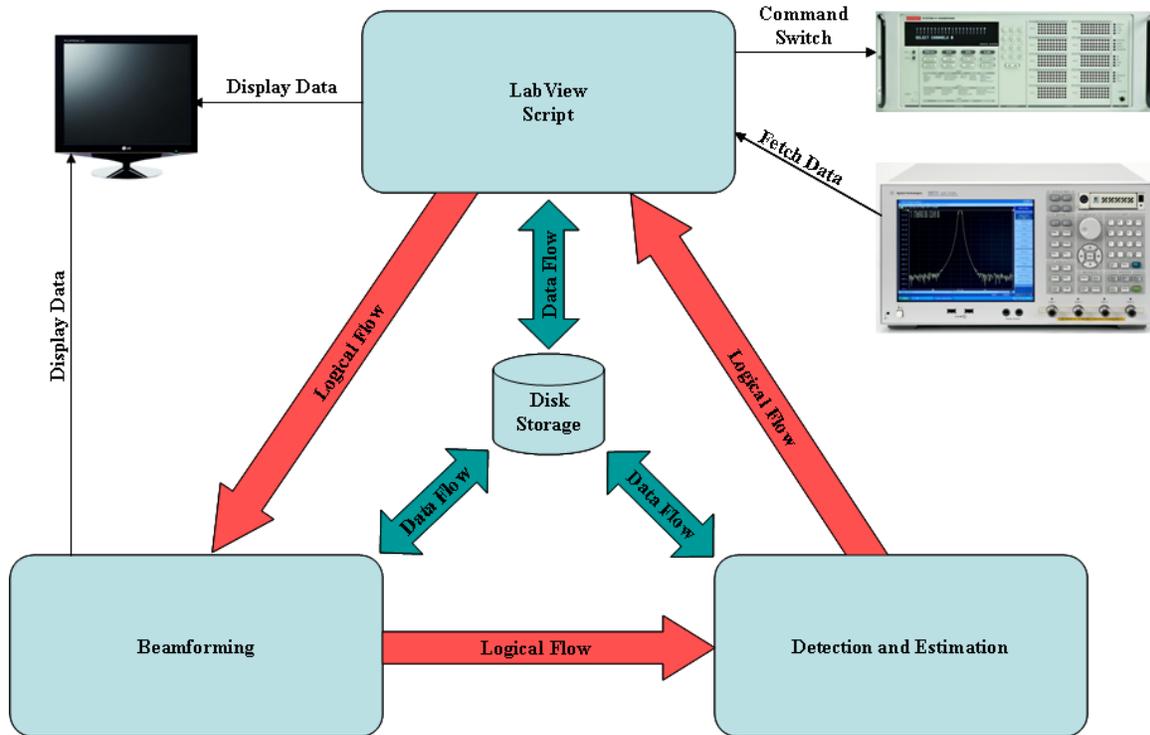


Figure 7. Software architecture for EV4 PARSEQ SAW RFID interrogator

4. Detection and Estimation

The remainder of this paper is devoted to a fairly detailed description of the design and performance of the signal detection and estimation algorithm incorporated in the current PARSEQ system. An introduction to the basic concepts of detection and estimation for SAW RFID signals measured using a network analyzer is presented in Appendix A. Throughout the remainder of the paper, we adopt the notation used in Appendix A to describe the signal processing techniques. In particular, we assume that the received signal to be processed for each of the pre-selected beams takes the form $\mathbf{R} = \mathbf{S} + \mathbf{N}$, where

$$\begin{aligned}\mathbf{R} &= (R_0, R_1, \dots, R_{N-1}), \\ \mathbf{S} &= (S_0, S_1, \dots, S_{N-1}), \\ \mathbf{N} &= (N_0, N_1, \dots, N_{N-1}).\end{aligned}$$

These three vectors represent the sampled frequency components for the total received signal in a particular beam, the received tag signal (which may be a composite of several tag response vectors), and the additive noise, respectively. The lowest sampled frequency, highest sampled frequency, and the number of sampled frequency components are all system configuration parameters. The noise samples are assumed to be complex-valued additive white Gaussian noise (AWGN) with mean zero and total variance σ^2 . As discussed in Appendix A, the variance is related to the power spectral density of the underlying real-valued continuous-time AWGN process observed at the network analyzer and the intermediate-frequency (IF) bandwidth setting on the analyzer.

The detection and estimation procedure described in Appendix A is a simple approach that works well in high signal-to-noise ratio (SNR) situations when there is little or no multiuser (i.e., multiple tag) or multipath interference within any beam. The approach implemented in the PARSEQ interrogator is considerably more sophisticated and performs much better in low SNR situations and in the presence of both multiuser and multipath interference. The implemented algorithm is a maximum-likelihood joint detection and estimation approach based on the wavelet transform. Although it is reasonably computationally intensive, it has been vectorized and optimized for fast computation using Matlab on a single processor. The algorithm can easily be partitioned for parallel implementation on multiple FPGA or DSP processors, and significant reductions in computation time can be realized in this manner.

4.1. Basic Approach

The algorithm operates by first computing a sampled version of the *continuous-time wavelet transform* of the so-called *pre-envelope* of the received signal with respect to the *wavelets* defined by the pre-envelopes of the RFID tag impulse response for each of the tags represented in the template data base. The wavelet transform [3] is precisely the right tool for our purposes because it provides a complete representation for an arbitrary time-domain signal in terms of time-shifted and time-scaled (i.e., compressed or dilated) versions of a particular template signal (the wavelet itself)¹. Since we know that changes in temperature cause SAW impulse-response waveforms to scale in time, and changes in range cause the waveforms to shift in time at the receiver, the problem of interest is joint signal detection and parameter estimation for arbitrarily time-scaled and time-shifted versions of known template signals in the presence of AWGN. It follows from this that maximization of the wavelet transform of the received signal with respect to the wavelets defined by the SAW templates in the data base leads to a maximum-likelihood procedure for detection and range/temperature estimation of SAW RFID tag signals [4].

We assume that all tags are bandlimited to the interval $[f_0, f_0 + B]$ Hz regardless of the temperature at which the tag is observed². If the tag frequency response is represented by the function $H(f)$, then the impulse response takes the form

$$h(t) = \int_{-\infty}^{\infty} H(f) e^{i2\pi ft} dt ,$$

and the sampled tag frequency-response vector will take the form $\mathbf{H} = (H_0, H_1, \dots, H_{N-1})$, where

$$H_k = H\left(f_0 + \frac{B}{N}k\right), \quad k = 0, 1, \dots, N-1.$$

The corresponding pre-envelope of the tag impulse response is then given by

¹ This is analogous to classical time-frequency representations such as the short-time Fourier transform and the Wigner-Ville transforms [2], but more mathematically rigorous.

² For our purposes, $f_0 \approx 2.36$ GHz and $B \approx 160$ MHz. The interval chosen for processing is actually a system configuration parameter, but all tags currently available satisfy this constraint

$$\tilde{h}(t) = \frac{2B}{N} \sum_{k=0}^{N-1} H_k e^{i2\pi(f_0 + \frac{B}{N}k)t}, \quad 0 \leq t \leq T = N/B,$$

and (ignoring any time-domain aliasing resulting from sampling in the frequency domain) we have the relationship

$$h(t) = \text{Re}[\tilde{h}(t)].$$

For the received signal vector $\mathbf{R} = (R_0, R_1, \dots, R_{N-1})$, we have the similar relationships:

$$R_k = R\left(f_0 + \frac{B}{N}k\right), \quad k = 0, 1, \dots, N-1,$$

$$r(t) = \int_{-\infty}^{\infty} R(f) e^{i2\pi ft} dt,$$

$$\tilde{r}(t) = \frac{2B}{N} \sum_{k=0}^{N-1} R_k e^{i2\pi(f_0 + \frac{B}{N}k)t}, \quad 0 \leq t \leq T = N/B,$$

$$r(t) = \text{Re}[\tilde{r}(t)].$$

The wavelet transform of the pre-envelope of the received signal with respect to the wavelet defined by the pre-envelope of the tag impulse response (assuming it has zero DC component) is given by

$$W_{r,h}(\alpha, \tau) = \frac{1}{\sqrt{\alpha}} \int_{-\infty}^{\infty} \tilde{r}(t) \overline{\tilde{h}\left(\frac{t-\tau}{\alpha}\right)} dt, \quad \forall \alpha > 0, \tau \in \mathfrak{R},$$

where the overbar indicates complex conjugation. If there are L tags in the template database with impulse-response signals $\{h_1, h_2, \dots, h_L\}$, and we know that exactly one tag is observed in a particular signal beam, then the maximum-likelihood estimate of the tag ID, temperature, and range is given by

$$\left(l^*, \alpha^*, \tau^*\right) = \underset{\{(l, \alpha, \tau) | l=1, \dots, L; \alpha > 0; \tau \in \mathfrak{R}\}}{\text{arg max}} \text{Re}\left[W_{h_l}(\alpha, \tau)\right].$$

To perform practical joint detection and estimation using this approach without knowing *a priori* how many tags are present, we use a technique called *successive interference cancellation* (SIC) [5], which is described below.

SIC, Stage 1. SIC is a two-stage joint detection and estimation process. To begin, we redefine the original pre-envelope received signal $\tilde{r}(t)$ as the residual signal for Stage 1 processing and compute a sampled version of $W_{r,h}$, given by

$$\mathbf{W}_{r,h}(i, j) = W_{r,h}(\alpha_i, \tau_j), \quad i = 1, 2, \dots, I; j = 1, 2, \dots, J,$$

for each of the impulse-response waveforms $h \in \{h_1, h_2, \dots, h_L\}$ in the template database. The grid of sampling points $\{(\alpha_i, \tau_j) \mid i = 1, 2, \dots, I; j = 1, 2, \dots, J\}$ in the wavelet domain is arbitrary and can be chosen during system configuration.

An example of this sampled wavelet transform is illustrated in Figure 8. For this example, the received signal is a time-scaled and time-shifted version of a single tag template embedded in AWGN at a signal-to-noise ratio (SNR) of 20 dB. In Part (a) of the figure, the chosen wavelet is also derived from the same template, while in Part (b), the wavelet was derived from a different tag template. In this case, the continuous wavelet transform was sampled at 800 time-shift points (approximately every 6.25 ns) and 41 time-scale points (approximately every 2.439° C). The time-scale and time-shift are referenced to the template itself; that is, the template is assumed to have been measured at a time shift of 0 and a time scale of 1. As the figure indicates, when the chosen wavelet matches the received signal, the magnitude of the wavelet transform shows a strong peak at the sampled time-scale, time-shift point closest to the true time-shift and time-scale of the received signal, which in this case was the point $(0.998182, 0.667128 \times 10^{-7})$. On the other hand, if the chosen wavelet does not match the received signal, the magnitude of the wavelet transform shows no strong peaks.

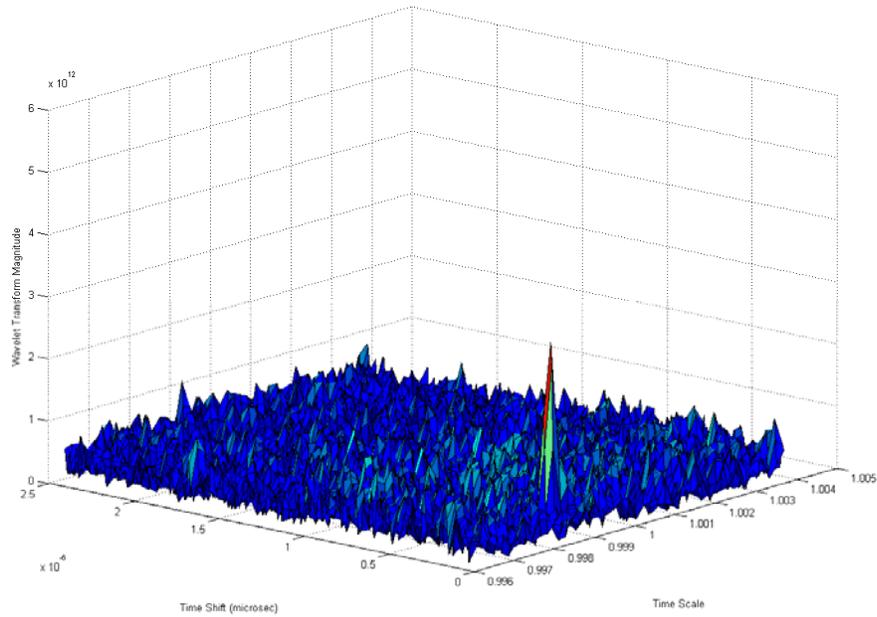
Using the entire collection of sampled wavelet transforms for all tag templates in the data base, we can solve a discretized version of the desired optimization problem:

$$(l^*, i^*, j^*) = \arg \max_{\{(l, i, j) \mid l=1, \dots, L; i=1, \dots, I; j=1, \dots, J\}} \operatorname{Re} \left[W_{r, h_l}(\alpha_i, \tau_j) \right],$$

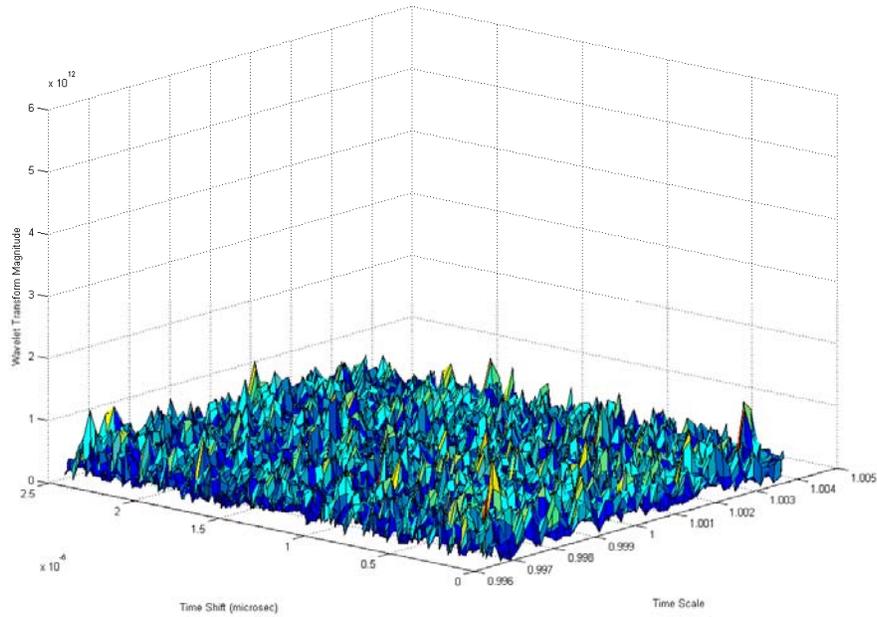
using simple exhaustive search. If $\operatorname{Re} \left[W_{r, h_{l^*}}(\alpha_{i^*}, \tau_{j^*}) \right]$ does not exceed a pre-selected detection threshold, we declare that no tags are present in the current beam. If $\operatorname{Re} \left[W_{r, h_{l^*}}(\alpha_{i^*}, \tau_{j^*}) \right]$ exceeds the detection threshold, we declare that some tag-like energy is present and perform secondary detection processing. Assuming tag l^* satisfies the secondary detection criteria as well, then we declare that tag l^* is present with the initial estimate of (α, τ) given by $(\alpha_{i^*}, \tau_{j^*})$. To compute the final estimate of (α, τ) for the detected tag, we solve the simplified continuous-time optimization problem

$$(\alpha^*, \tau^*) = \arg \max_{\{(\alpha, \tau) \mid \alpha > 0; \tau \in \mathbb{R}\}} \operatorname{Re} \left[W_{h_{l^*}}(\alpha, \tau) \right]$$

using a numerical optimization routine with starting value $(\alpha_{i^*}, \tau_{j^*})$. The numerical optimization routine is the most computationally complex part of the algorithm, with a complexity of $\mathcal{O}(N^3)$.



(a) Wavelet matched to signal.



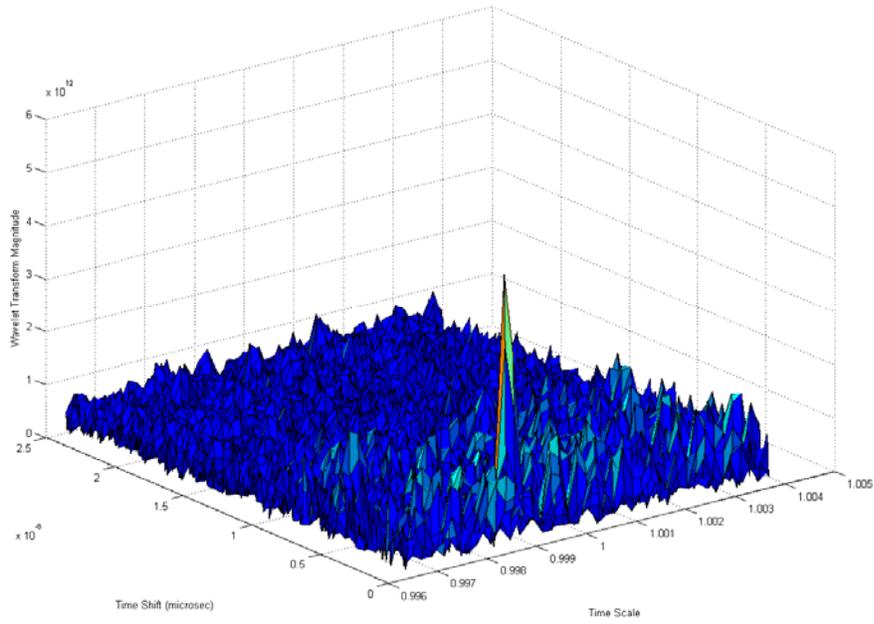
(b) Wavelet not matched to signal

Figure 8. Sampled wavelet transform for received signal representing single tag response at SNR of 20 dB.

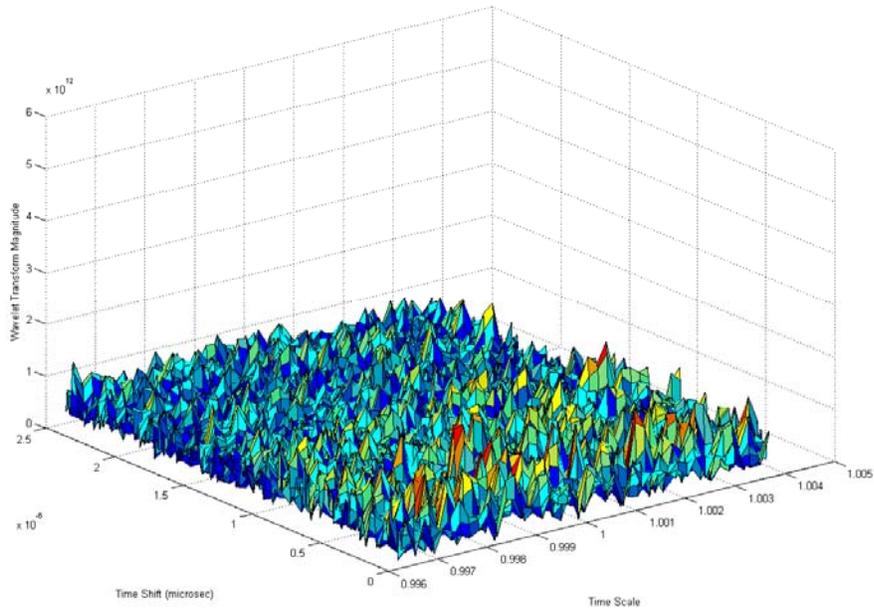
SIC, Stage 2. If no tags are detected in the current residual signal during Stage 1, then detection and estimation processing for the current beam is terminated and processing proceeds to the next set of beam data. If a tag is detected in the current residual signal, then for the second stage of

processing, we assume that the estimated tag ID, temperature, and range from Stage 1 are all correct and reconstruct an estimated waveform corresponding to that particular tag signal by projecting the current residual signal onto the appropriately time-scaled and time-shifted version of the detected tag template from the data base. This reconstructed signal is then treated as a known interfering signal in the current residual signal, and a new residual signal is computed by subtracting the reconstructed signal from the current residual. This completes Stage 2 of the SIC processing, and control is passed back to Stage 1 to search for additional tags in the current beam.

An example of the sampled wavelet transform before and after SIC processing is illustrated in Figure 9. For this example, the received signal is a superposition of time-scaled and time-shifted versions of two different tag templates embedded in AWGN. The SNR for each of the individual tag response signals, ignoring the interference due to the other tag, is once again set at 20 dB . The figure was computed using the wavelet defined by the tag template that produced the maximum observed sampled wavelet magnitude among the 11 templates in the data base. As expected, this was one of the two tag templates represented in the composite received signal. As before, the continuous wavelet transform was sampled at 800 time-shift points and 41 time-scale points, with the time-scale and time-shift referenced to the template defining the wavelet. Part (a) of the figure shows the sampled wavelet transform of the original composite signal during the first iteration of Stage 1 processing prior to any SIC processing. Part (b) of the figure shows the wavelet transform of the residual signal in the second iteration of Stage 1, after SIC processing during the first iteration of Stage 2 to remove the signal component corresponding to the first detected signal, which in this case was the signal matched to the chosen wavelet. As the figure indicates, prior to SIC processing the magnitude of the wavelet transform shows a strong peak at the sampled time-scale, time-shift point closest to the true time-shift and time-scale of the detected signal, which in this case was the point $(0.998545, 0.800554 \times 10^{-7})$. After SIC processing, the magnitude of the wavelet transform shows no strong peaks, indicating that the signal component corresponding to the detected signal was successfully removed.



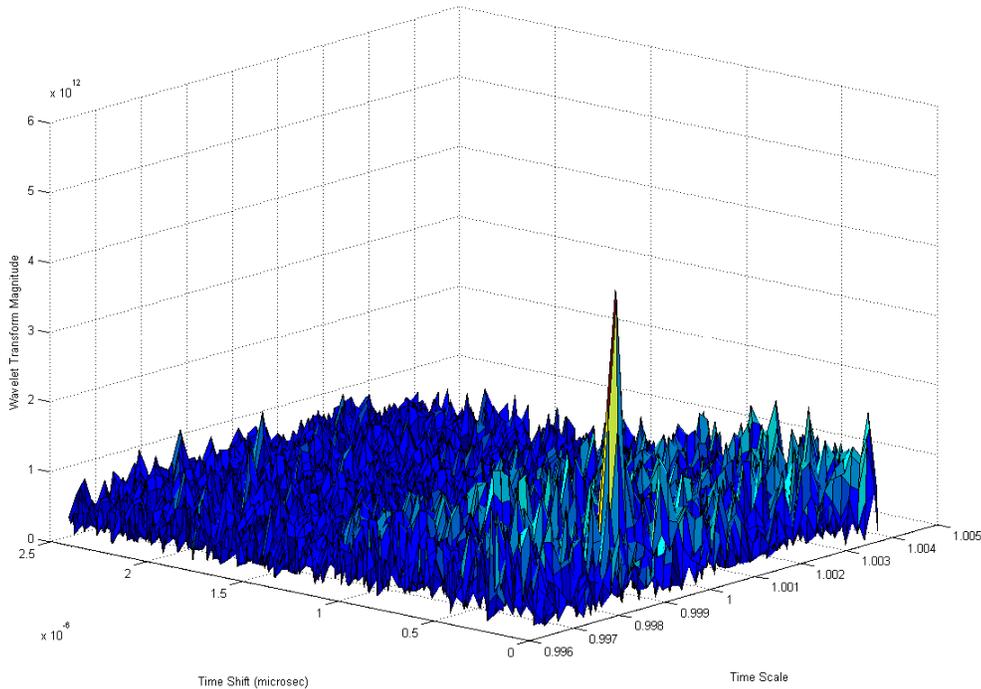
(a) Wavelet transform prior to SIC processing (first iteration of Stage 1).



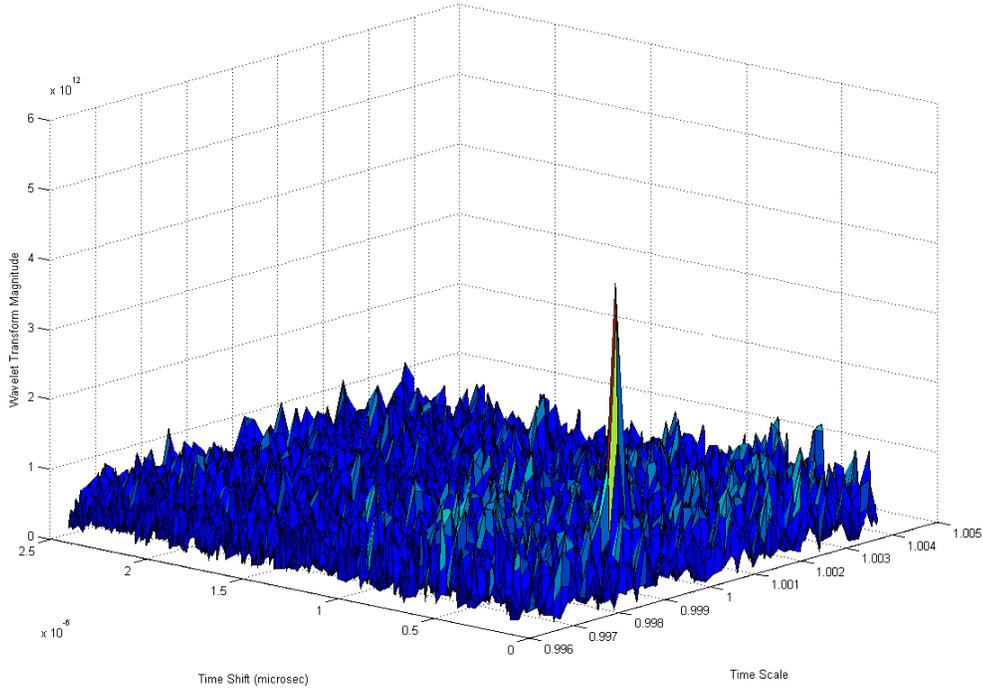
(b) Wavelet transform after SIC processing (second iteration of Stage 1).

Figure 9. Sampled wavelet transform for received signal representing superposition of two tag responses at SNR of 20 dB using wavelet with strongest response (detected signal) during first iteration of Stage 1.

For reference, we also include Figure 10, which again shows the sampled wavelet transform of the composite signal before and after SIC processing, but in this case, the figure was computed using the wavelet matched to the second of the two tag responses represented in the composite signal, which was the wavelet with the second strongest response in the first iteration of Stage 1. Part (a) shows the wavelet response during the first iteration of Stage 1 prior to any SIC processing to remove the first signal detected in Stage 1. The maximum magnitude of this response is slightly less than the maximum magnitude of the wavelet response illustrated in Figure 9(a), indicating that this tag response was not detected during the first iteration of Stage 1 processing. Part (b) of Figure 10 shows the wavelet transform of the residual signal in the second iteration of Stage 1, after SIC processing during the first iteration of Stage 2 to remove the first detected signal component. This response was in fact the strongest response among all the remaining templates in the data base (excluding the template for the first detected signal) during the second stage of Stage 1 processing. Hence, the second tag response represented in the composite received signal was detected during the second iteration of Stage 1. As the figure indicates, the magnitude of the wavelet transform still shows a strong peak at the sampled time-scale, time-shift point closest to the true time-shift and time-scale of the detected signal, which in this case was the point $(0.998182, 0.667128 \times 10^{-7})$. This indicates that very little of this signal component was removed during the first stage of SIC processing, which is the desired result. For this example, processing terminated during the third iteration of Stage 1, indicating that no further tag responses could be detected.



(a) Wavelet transform prior to SIC processing (first iteration of Stage 1).



(b) Wavelet transform after SIC processing (second iteration of Stage 1).

Figure 10. Sampled wavelet transform for received signal representing superposition of two tag responses at SNR of 20 dB using wavelet with second strongest response during first iteration of Stage 1.

4. Performance Evaluation

In this section, we present an evaluation of the detection and estimation performance for the PARSEQ system. The performance of the system was evaluated both experimentally and using Monte Carlo computer simulations. The experimental performance was evaluated for the system as a whole, including the effects of the beamformer, but only estimation performance (as opposed to both detection and estimation performance) was quantified. The simulated performance was evaluated only for the detection and estimation algorithm itself, but both detection and estimation results were quantified. The relative efficiency of the estimation algorithm was evaluated with respect to the Cramer-Rao (C-R) lower bounds on the variance of unbiased estimates of both temperature and range. Analytical expressions for the (C-R) bounds were derived and evaluated numerically for comparison to the simulated performance results.

4.1. Monte Carlo Performance Evaluation

For the Monte Carlo performance evaluation of the detection and estimation algorithm, the tag template data base consisted of 11 tag waveforms that were measured experimentally by connecting each of the 11 reference tags to the S1 port of the Agilent ENA network analyzer with an RF interconnect cable and recording the frequency response from each tag at 801 equally spaced points in the 160 MHz interval from 2.36 GHz to 2.52 GHz. The resulting measured frequency-response vectors were then processed to remove out-of-band noise, the two synch pulses, and other non-tag-related artifacts using the following procedure:

1. The tag response was filtered to the desired bandwidth by multiplying the measured frequency-response vector by a 801-point Blackman frequency-domain window.
2. The filtered frequency-response vector was converted into an over-sampled base-band impulse-response vector by padding the filtered frequency-response vector with zeros to a length of 65,536 samples and then performing an inverse FFT to produce a 65,536-point base-band impulse-response vector.
3. The base-band impulse-response vector was windowed to remove the synch pulses and other non-tag-related transients by multiplying with a rectangular, unit-amplitude time-domain window with non-zero response in the interval from 1.26 μs to 2.63 μs and then normalized to have an average single-sample power of unity.
4. The normalized impulse-response vector was converted back into the frequency domain by performing a 65,536-point FFT and then truncated back to 801 points to produce the final normalized frequency-response vector to be stored in the template database.

The templates were all collected at room temperature, and the precise temperature of each tag during the measurement process was not recorded. Instead, an arbitrary reference temperature of 20° C was assigned to each tag template for purposes of performance evaluation. Since the physical quantity being simulated and measured in this performance evaluation was actually time dilation of the tag waveform, from which temperature change was then inferred, the true temperature of each tag during the template measurement procedure was not relevant. An example of one of the resulting tag template impulse-response waveforms is illustrated in Figure 11.

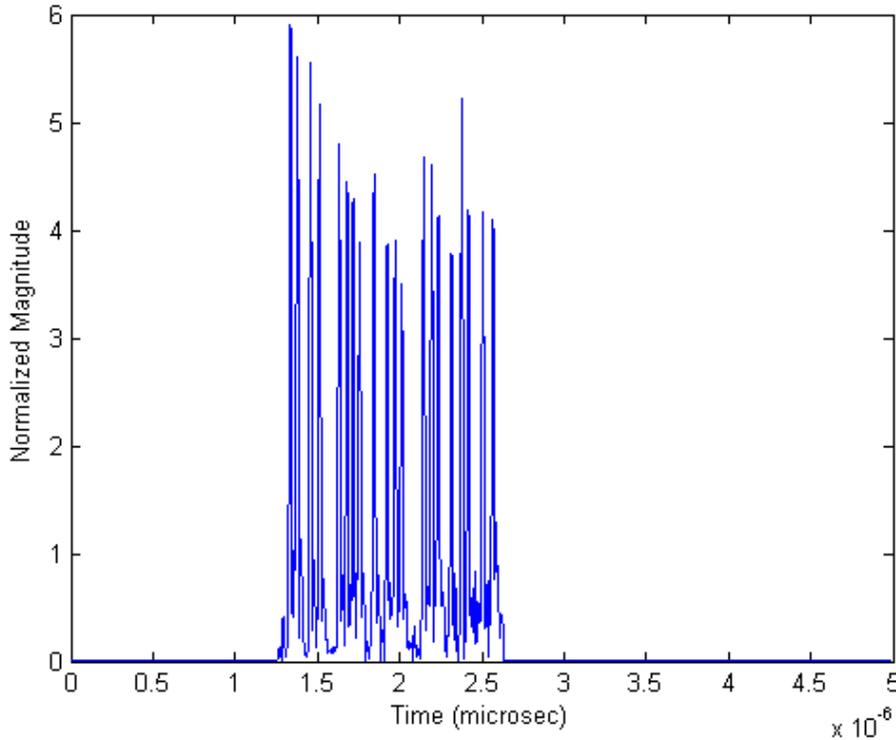


Figure 11. Representative tag template impulse response.

To simulate the detection and estimation performance at a range of different temperatures and signal-to-noise ratios, each of the tags was “observed” at a different temperature and at many different SNRs by first numerically dilating or compressing the template impulse-response waveform to the scale corresponding to the temperature assigned to that tag for the simulation³ and then adding simulated complex-valued, circularly symmetric, additive white Gaussian noise (AWGN) to the resulting tag frequency response vector. In particular, exactly one of the temperatures from the set $\{-5, 0, 5, 10, 15, 20, 25, 30, 35, 40, 45\}$ was assigned to each of the 11 tags in the template database and all tags were simulated at 16 different SNR levels chosen at 1 dB increments in the range from 15 dB to 30 dB. For all simulations, the detection thresholds in the algorithm were fixed at levels selected experimentally to give good performance at a nominal SNR of 20 dB. Performance results for each SNR tested could have been improved by optimizing the detection thresholds for that SNR; however, such SNR-specific tuning is often not a realistic option for actual applications, so we used the more realistic approach of picking detection thresholds for a fixed nominal SNR value.

4.1.1. Simulation Results

For this set of simulations, a simulated received signal was derived by superimposing the simulated temperature-scaled frequency response vectors for multiple tags and then adding AWGN to produce the desired SNR level. That is, the SNR level was the same for all tags in the composite signal and was computed without consideration of the mutual interference produced by the multiplicity of tags. Detection and estimation results were simulated for composite received signals containing only a single tag, two tags, etc., up to the full collection of 11 tags represented in the template database. The full range of SNR levels from 15 to 30 dB was simulated for all tag multiplicity levels.

The detection performance results for the Monte Carlo simulations are presented in Figures 12 and 13 below. Figure 12 shows the probability of detection (P_d) for the full range of SNR values tested as a function of the number of total tag response signals present in the received signal. Here, P_d was computed by counting the number of tags in each simulated composite received signal that were correctly detected and dividing by the total number of tag response signals included in the composite. Notice that for the lowest SNR level of 15 dB, the detection probability was quite high (93.8%) when only a single tag was present but quite low (34.9%) when all 11 tags were present. However, as the SNR increases, detection probability grows monotonically regardless of the number of tags present, and for SNR levels greater than or equal to the nominal value of 20 dB, detection probability exceeds 94% uniformly across the entire range of tag multiplicity. In fact, at all SNR levels above 21 dB, detection probability uniformly exceeds 99%.

³ For the simulation, tag TCD was assumed to vary linearly at the rate of $7.285714285714435 \times 10^{-5}$ per degree C.

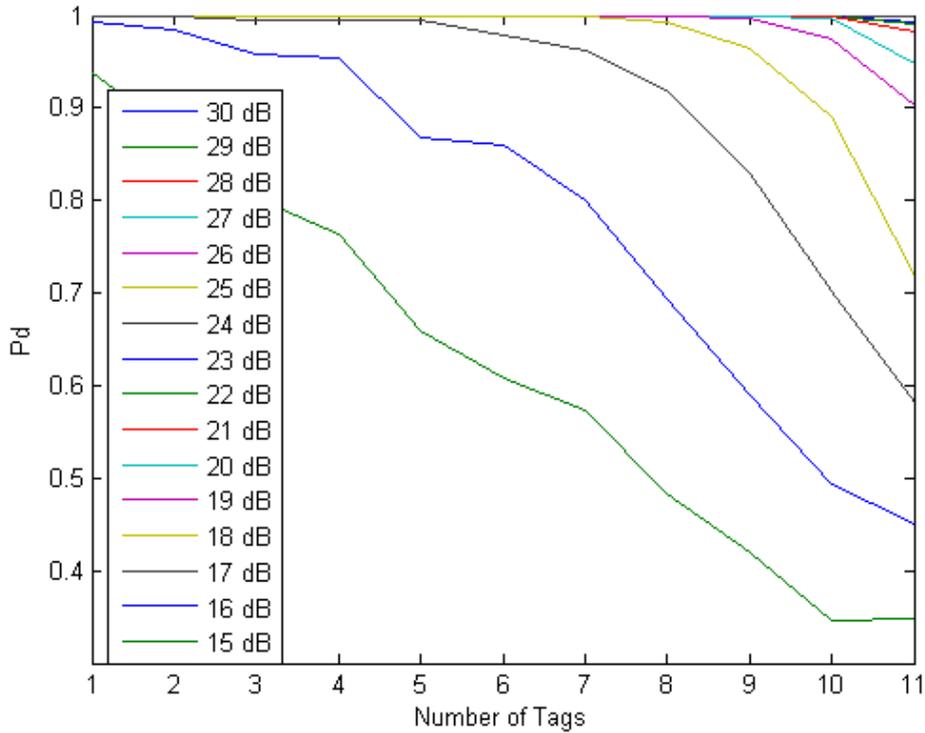


Figure 12. Probability of detection as a function of SNR and tag multiplicity.

Figure 13 shows the probability of false alarm (P_{fa}) or *false-alarm rate* for the simulations. Here, P_{fs} was computed by counting the number of tags in each simulated composite received signal that were reported as detections but not actually present in the signal and dividing by the total number of tag templates *not* present in the composite. Notice that the false-alarm rate is less than 5% uniformly across the range of tag multiplicity for all SNR levels below 26 dB, and that it is less than 1% uniformly across all SNR levels as long as the number of tags present in the signal is less than 5. However, for SNR levels above 25 dB, the false-alarm rate begins to grow rapidly as the number of tags present in the received signal increases above 5, and for very high SNR levels, quickly reaches an unacceptable level. This is a result of the fact that some extraneous tag signal components always remain in the residual signal after the completion of the SIC procedure described above, and at sufficiently high SNR levels, the correlation between such extraneous components and the remaining tag templates can produce peaks in the wavelet transform that cannot be distinguished from true signal peaks. This problem can be ameliorated to some extent by introducing some adaptation into the selection of detection thresholds, but we are still experimenting with that procedure.

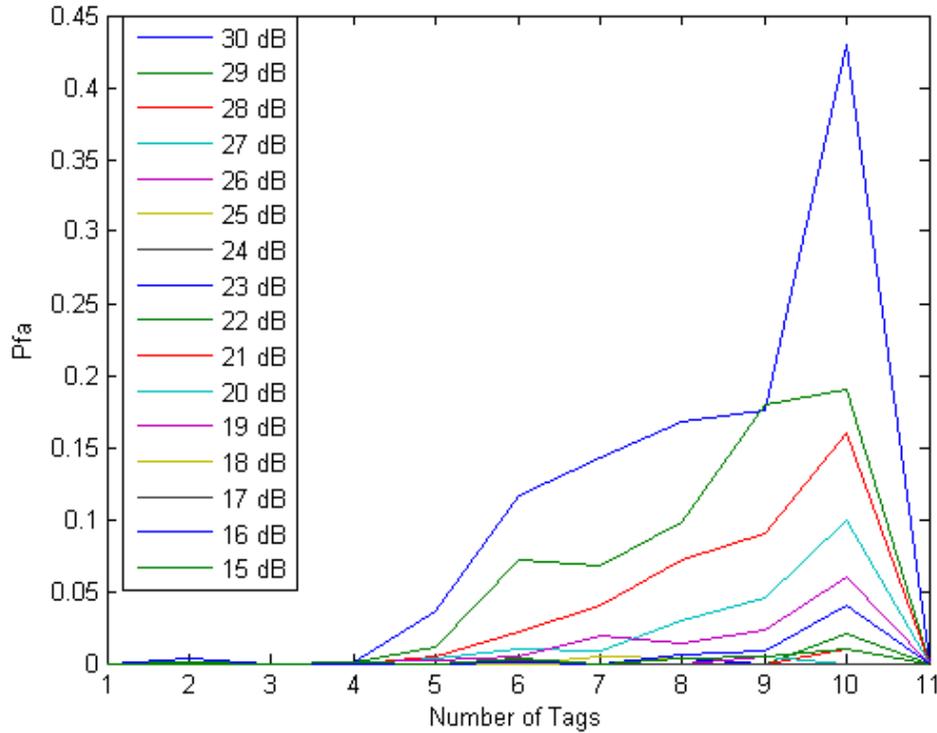


Figure 13. Probability of false alarm as a function of SNR and tag multiplicity.

The estimation performance results for the Monte Carlo simulations are presented in Figures 14-17 below. Figures 14 and 15 show the mean squared error (MSE) for estimation of temperature and range, respectively, for the full range of SNR values tested as a function of the number of total tag response signals present in the received signal. As discussed in Section 4.1, both quantities were estimated simultaneously, but statistical analysis of the empirical results indicate that the two estimates are very weakly correlated in the low MSE regime. Also, while the evidence suggests that at least one of the estimators is not truly unbiased, which is discussed further below, the observed squared bias for both estimators was always at least an order of magnitude less than the observed MSE in the low MSE regime. Hence, the contribution of the bias to the empirical MSE of the estimates is insignificant in all scenarios where the temperature and range estimators perform well, so only the total MSE is presented here as a performance criterion.

It should also be noted that estimation performance was tabulated only for correctly detected tags. That is, if a tag was not reported as detected, then no attempt was made to estimate either temperature or range, so it was automatically excluded from the statistical analysis. Similarly, if a tag was falsely reported as detected, and estimates of temperature and range were both computed, those estimates were excluded from the statistical analysis of estimation performance.

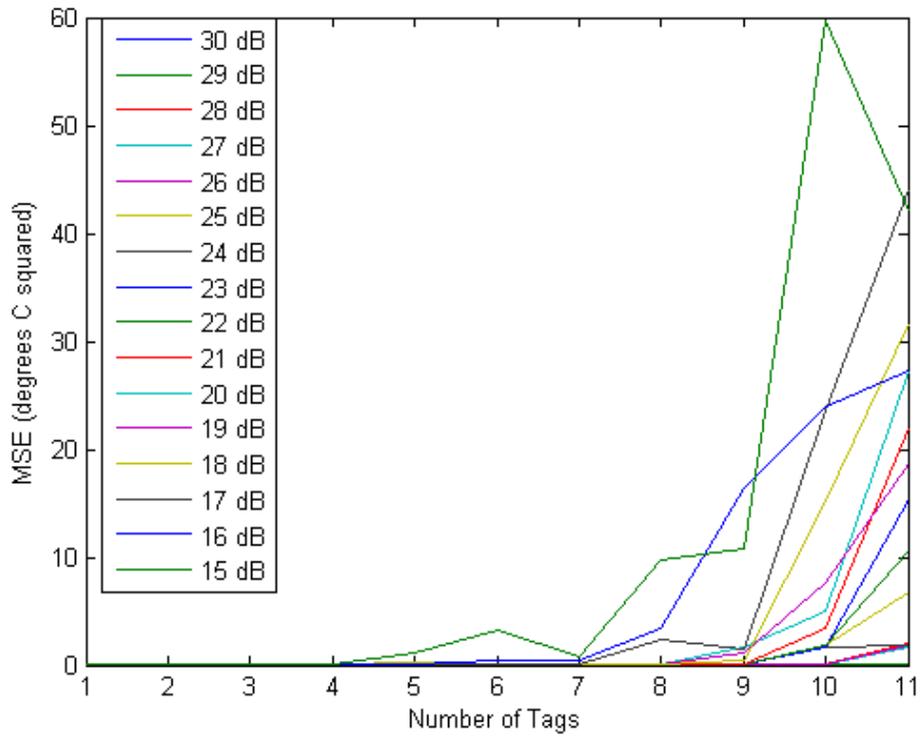


Figure 14. MSE for temperature estimation as a function of SNR and tag multiplicity.

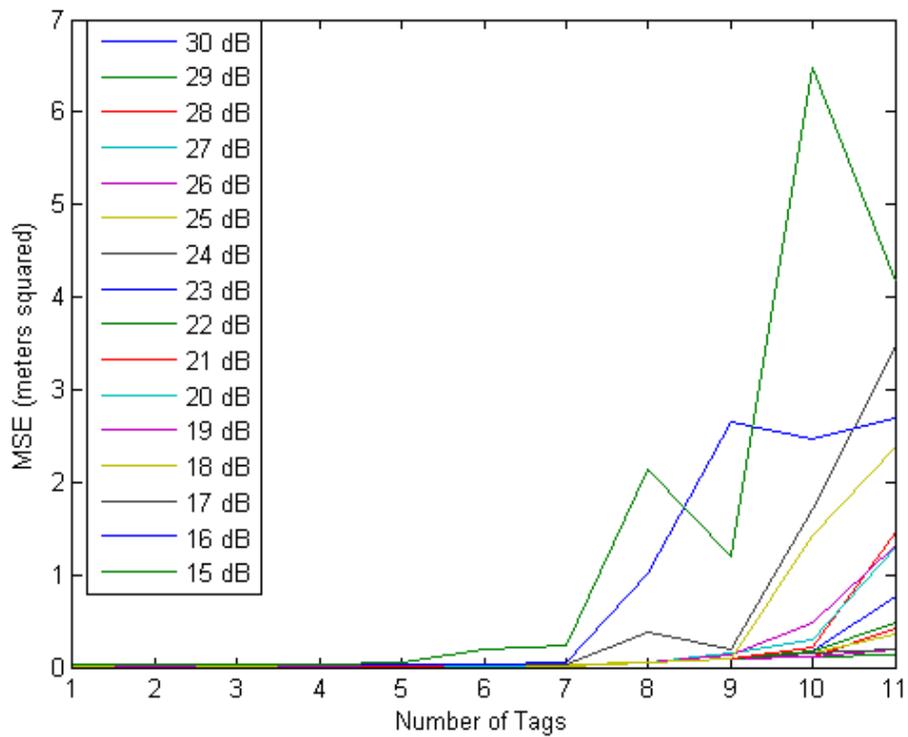


Figure 15. MSE for range estimation as a function of SNR and tag multiplicity.

Notice that both figures indicate very good estimation performance across the entire range of SNR levels until the number of tags in the received signal exceeds 7. In fact, for all SNR levels in excess of 15 and tag multiplicities less than or equal to 7, the MSE for temperature is a small fraction of 1° C and the MSE for range is a small fraction of 1 meter. Notice also that for each SNR level below 26 dB, there is a very rapid increase in MSE for both temperature and range when the number of tags in the signal exceeds a certain level. Extrapolating just a bit, this would seem to indicate that there is a clear performance breakdown threshold for each SNR level, such that for tag multiplicities below that threshold, estimation performance is very good, and for tag multiplicities above that threshold, estimation performance quickly becomes very bad. For the lowest SNR tested (15 dB), this threshold seems to be 7 tags, while for SNR levels above 25 dB, the threshold seems to exceed 11.

We now consider the estimation performance in the absence of any tag mutual interference; that is, when only a single tag is present in the received signal and estimation performance is entirely a function of SNR. These results are presented in Figures 16 and 17, where the MSE for temperature and range are plotted, respectively, in comparison to the Cramer-Rao (C-R) lower bound on the variance (equivalently, MSE) of unbiased estimates for both quantities [6].

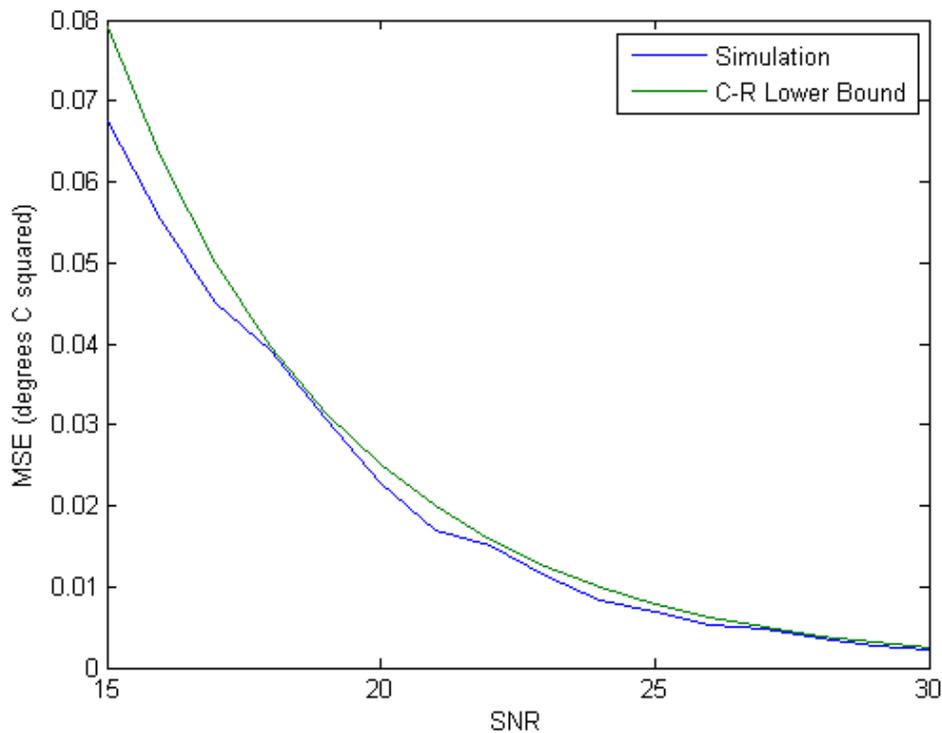


Figure 16. MSE for temperature estimation as a function of SNR for a single tag response.

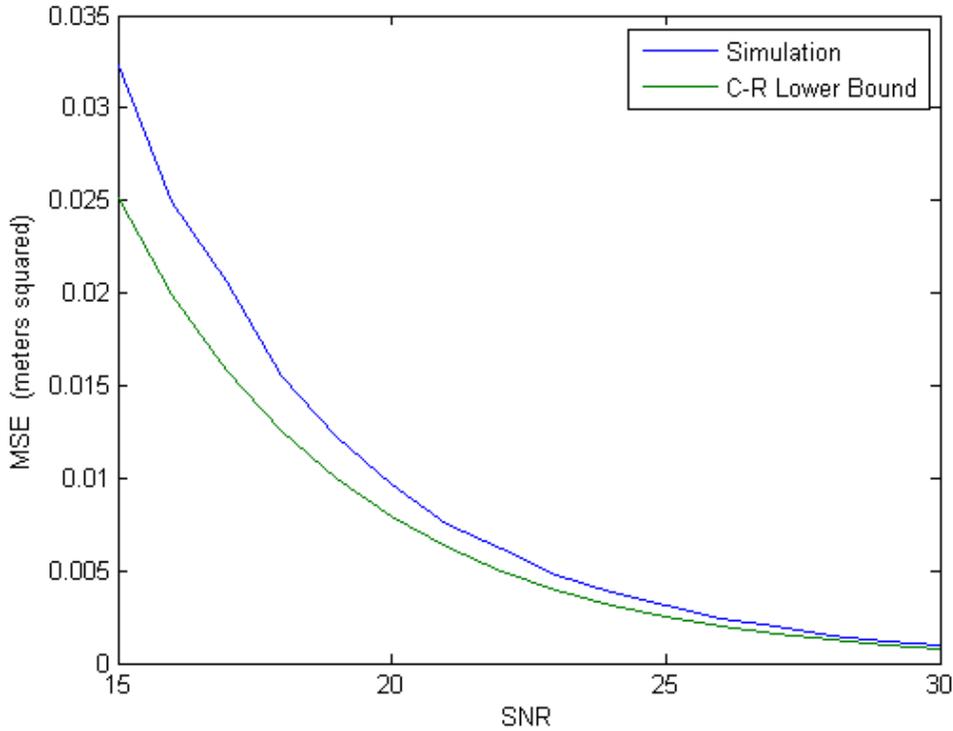


Figure 17. MSE for range estimation as a function of SNR for a single tag response.

Two conclusions can be drawn from the results presented in Figures 16 and 17, namely, the estimates of temperature and range in the single-tag scenario are both very accurate and very efficient; however, some discussion is in order to clarify this. First, the actual MSE for both temperature and range is underestimated by the empirical MSE values in both figures at SNR levels 15 and 16. This is a result of the fact that not all tags in the simulated data set were correctly detected at these SNR levels. Although only a few tags were missed in each case, they were necessarily the weaker tags, so the average SNR of the detected tags was actually somewhat higher than the specified level. This implies that the empirical estimates of MSE in these two cases were necessarily biased below the true MSE of the estimators. If estimates of temperature and range for the missing tags had been computed and included in the statistical analysis, the observed values of MSE would have risen slightly.

Second, the fact that the empirical MSE curve in Figure 16 actually falls *below* the C-R lower bound curve is not an error. While, the C-R bound is a true lower bound on the MSE for any *unbiased* estimator of a parameter, it is possible in some cases for a *biased* estimator to have MSE lower than the C-R bound. Hence, the results in Figure 16 indicate that the bias of the temperature estimator studied here is not zero even though, as discussed above, it was observed to make an insignificant contribution to the error in the low MSE regime. Nevertheless, the MSE of the estimator is actually smaller than the MSE of an optimal (minimum MSE or MMSE) unbiased estimator, so the temperature estimator can be regarded as “*super-efficient*” relative to the MMSE unbiased temperature estimator.

Finally, the C-R lower bound used for range estimation in Figure 17 is probably not the tightest possible lower bound for range estimation when temperature and range are estimated jointly,

since it actually represents the tightest possible C-R bound for estimation of range under the assumption that temperature is known a-priori. (i.e., temperature and range *not* estimated jointly). Even using this loose lower bound, the empirical MSE curve for the range estimates is uniformly less than 30% greater than the C-R lower bound curve across all SNR levels and approaches 20% at the highest SNR levels. Alternatively, shifting the empirical MSE curve to the left by 1 dB would put the curve uniformly below the C-R lower bound curve. Hence, while the range estimator does not outperform the MMSE unbiased range estimator, the SNR penalty is less than 1 dB relative to MMSE unbiased estimation of range alone.

4.2. Experimental Performance Evaluation

The experimental performance analyses performed on the PARSEQ system to date allow only a qualitative evaluation of the estimation performance of the system. They are included here in order to at least partially validate the simulation results as well as to convey some information about the estimation performance of the system in a dynamic environment. The discussion of the experimental performance results presented in this section is taken largely from an earlier technical report prepared by Mr. Reeve Ingle [7]. Minor editorial changes have been made here by the author.

4.2.1. Experiment Description

A single SAW RFID tag was placed in a temperature-controlled thermal chamber. An eight-foot length of M17/158 coaxial cable was passed through an insulated port in the thermal chamber and used to connect the RFID tag inside the chamber to a single antenna element mounted on a tripod outside of the thermal chamber. This antenna was placed at a distance of two meters bore-site to the 72-element receiver antenna array. The transmit antenna was mounted on a tripod placed beside the receiver antenna array and pointed at the single-element antenna connected to the tag. The network analyzer IF bandwidth was set to 15 kHz during all thermal tests.

The thermal chamber temperature was allowed to settle at 20° C to begin the experiment. A 15-minute thermal soak was performed at 20° C to allow the RFID tag to stabilize at this temperature. At this point, calibration parameters were set in the software to calibrate the system to generate a temperature reading of 20° C. Next, the thermal chamber temperature was decreased to -73° C at a constant temperature ramp rate (usually 2.5 °C/min). A thermal soak was again performed at -73° C for a minimum of 15 minutes to allow the tag to stabilize at that temperature. Subsequently, the thermal chamber temperature was increased back to 20° C at a constant temperature ramp rate (usually 4.5 °C/min). Throughout trials of this experiment, thermal chamber temperature data were time-stamped and recorded to an output file, and RFID tag range and temperature estimates were time-stamped and recorded to an output file.

4.2.2. Results & Analysis

As the thermal chamber temperature was decreased from 20° C to -73° C, the chamber was unable to maintain a constant temperature ramp rate over the entire temperature range. The temperature ramp rate decreased significantly as the chamber temperature approached -73° C, which is near the lower limit of temperature achievable with this chamber. Fortunately, the chamber was able to maintain a constant temperature ramp rate when the temperature was increased from -73° C to 20° C. Thus, estimation results are presented only for tests in which the temperature was increased from -73° C to 20° C. Also, due to a minor software bug that was later corrected, range measurements exhibited a temperature-dependent bias with a downward

trend as temperature increased. The actual range of the tag antenna was fixed at 2 meters throughout all trials.

4.2.2.1. Temperature Results

Tests were performed at several different temperature ramp rates to test the ability of the PARSEQ system to track a time-varying temperature accurately. The results of these tests are illustrated in Figures 18-20 below.

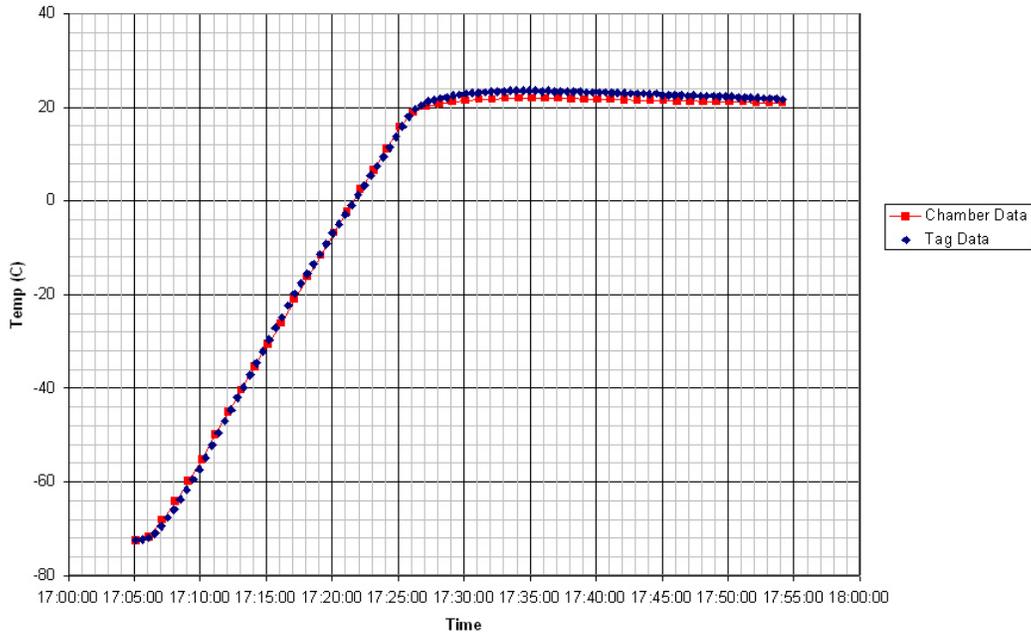


Figure 18. Comparison of measured temperature to actual chamber temperature for a temperature ramp rate of 4.5 °C/min.

Figure 18 displays a plot of the actual thermal chamber temperature (red trace) and the temperature measurements from the RFID tag (blue trace) as the temperature was increased from -73° C to 20° C at a temperature ramp rate of 4.5 °C/min. As the figure indicates, the RFID temperature measurements began approximately 0.2° C above the actual chamber temperature and ended approximately 1° C above the actual chamber temperature after a 25-minute thermal soak at 20° C. Note that the system tracked the actual temperature measured inside the chamber to within an accuracy of less than 1° C over the entire ramp but overshoot slightly at the end of the ramp.

It should be noted that switching between all 72 elements in the receiver antenna array requires approximately 20 seconds, and the thermal chamber temperature increases by approximately 1.5° C during that time at a ramp rate of 4.5 °C/min. Thus, each antenna element receives a tag waveform that is scaled differently than the tag waveform received by the previous element, and the output of the beam-forming program is “smeared out” in time. The output temperature computed by the detection and estimation algorithm based on the beam-formed data seems to equal roughly the average chamber temperature during the 72-element sweep. During subsequent trials, the temperature ramp rate was increased to try to determine the maximum rate that the system is capable of tracking. That is, above a certain temperature ramp rate, it is to be expected

that the beam-formed signal will be so badly distorted that the estimation algorithm will no longer be able to either detect the tag or produce an accurate temperature estimate.

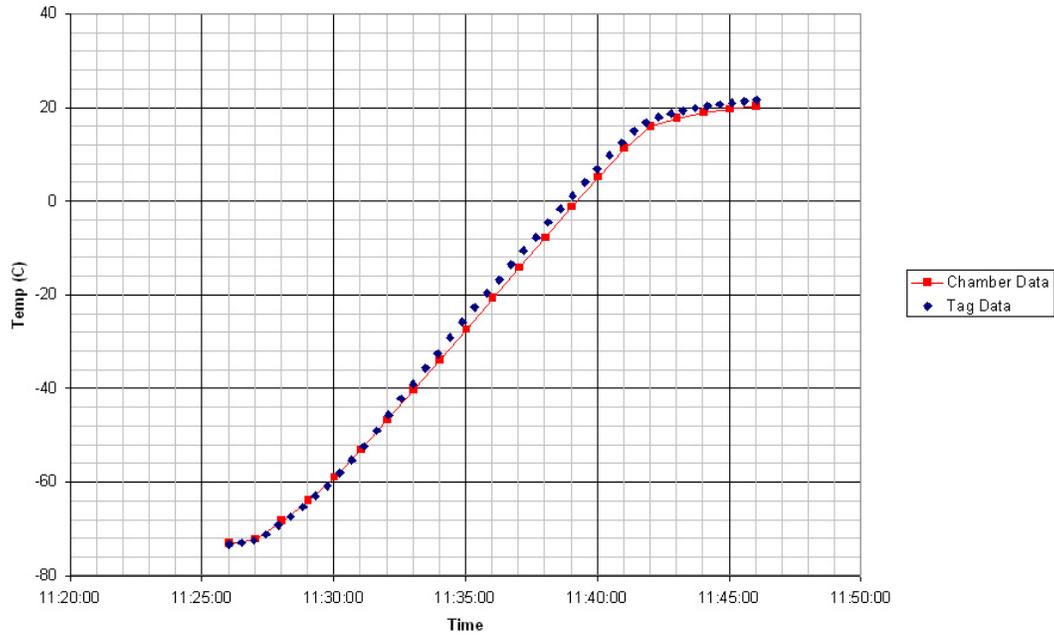


Figure 19. Comparison of measured temperature to actual chamber temperature for a temperature ramp rate of 6 °C/min.

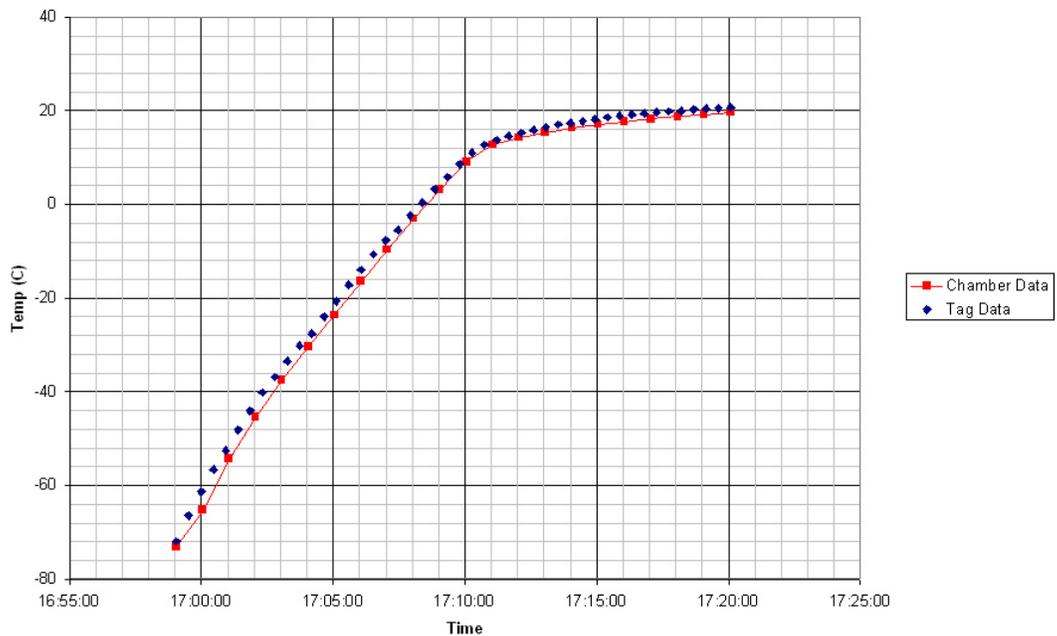


Figure 20. Comparison of measured temperature to actual chamber temperature for a temperature ramp rate of 11 °C/min.

Figure 19 shows the thermal chamber temperatures and RFID tag temperatures after the temperature ramp rate was increased to 6 °C/min. At this rate, the thermal chamber temperature

changes 2° C during the time it takes to sweep all 72 antenna elements. The RFID temperature measurements generally overstated the actual temperature in the chamber slightly but remained within approximately 2° C of the actual thermal chamber temperature throughout the test. Similarly, Figure 20 shows the results of the experiment when the temperature inside the chamber was increased at the maximum possible rate, which was approximately 11 °C/min at the peak. At this rate, the thermal chamber temperature changes approximately 3.7° C during the time it takes to sweep all 72 antenna elements. Once again, the temperature measurements overstated the actual chamber temperature slightly but remained within approximately 2° C of the actual temperature throughout the test.

Hence, although it is clear that the system will fail to track accurately at some rate of temperature change, the PARSEQ system was able to track temperature reasonably accurately at all ramp rates tested in these experiments.

4.2.2.2. Range Results

An example of one set of range measurement results is illustrated in Figure 21 below. These measurements were taken as the chamber temperature was increased from -73° C to 20° C at a ramp rate of 4.5 °C/min. As discussed above, a software bug that has since been corrected created a temperature-dependent bias in the range estimates that makes it difficult to draw many conclusions from these results. Nevertheless, if the linear trend is removed from the displayed data, it is clear that the variability of the range estimates falls within a range of approximately ±0.2 meters over the entire temperature ramp.

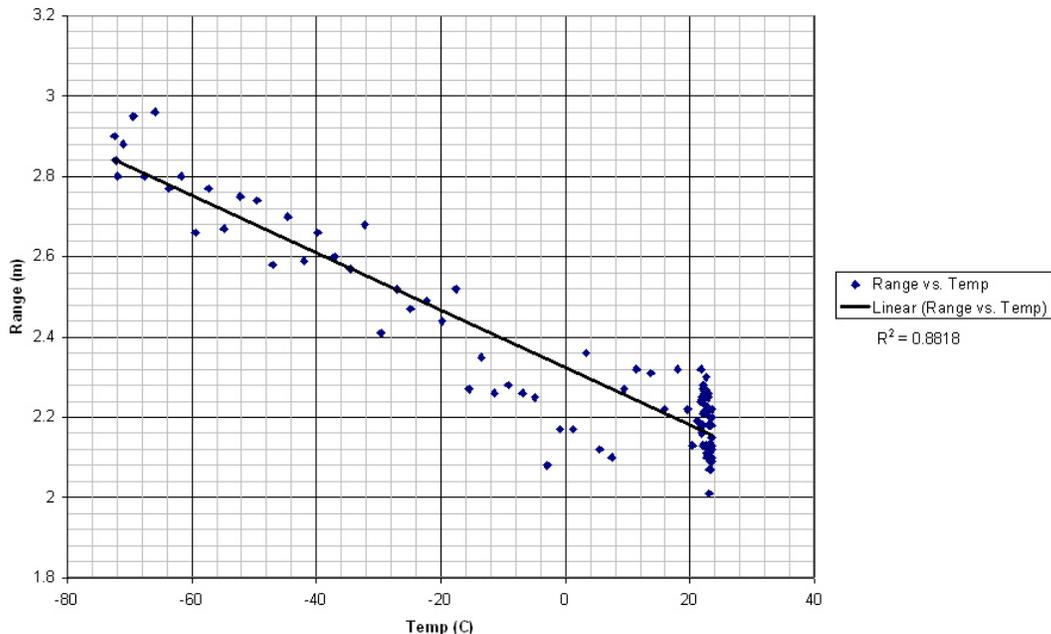


Figure 21. Plot of range estimates vs. temperature for a temperature ramp rate of 4.5 °C/min and fixed actual range of 2 m.

5. Conclusion

This paper briefly summarizes the design of the NASA JSC PARSEQ long-range SAW RFID interrogation system and presents a detailed description of the design and performance of the

signal detection and estimation algorithms incorporated in the system. Numerical performance results indicate that the detection and estimation algorithm can detect and discriminate between at least 11 superimposed tag waveforms with very high probability in composite signals corrupted with AWGN at SNR levels as low as 20 dB. At the same SNR levels, the algorithm produces extremely accurate estimates of both temperature and range for up to 8 tags simultaneously, and continues to produce reasonably accurate estimates with as many as 11 tags present simultaneously. For all SNR levels above 25 dB, temperature and range estimates are extremely accurate for all tag multiplicity levels tested. Further, for signals containing only a single tag, the temperature estimation outperforms the best possible unbiased estimator (in terms of MSE) for the entire range of SNR levels tested, while the range estimation performance comes within 30% of that of the best possible unbiased estimator over the same SNR range.

In terms of experimental performance, the available results qualitatively validate the numerical performance results and indicate that the PARSEQ system can reliably estimate temperature and range in real time and track temperature variation at rates up to 11 °C/min.

6. References

- [1] Hartmann, C. S. and Claiborne, L. T., “Fundamental Limitations on Reading Range of Passive IC-Based RFID and SAW-Based RFID,” *Proceedings of the 2007 IEEE International Conference on RFID*, pp. 41-48, Dallas, TX, March 26-28, 2007.
- [2] Hartmann, C. S., Brown, P., Bellamy, J., “Design of Global SAW RFID Tag Devices,” *Proceedings of the Second International Symposium on Acoustic Wave Devices for Future Mobile Communication Systems*, Chiba University, Japan, March 2004.
- [3] Daubechies, I., *Ten Lectures on Wavelets*, SIAM, Philadelphia, 1992.
- [4] Ostermeyer, G., Pohl, A., Hausleitner, C., Reindl, L., Seifert, F., “CDMA for Wireless SAW Sensor Applications,” *Proceedings of the 4th International Symposium on Spread Spectrum Techniques and Applications*, vol. 2, pp. 795-799, Sept. 22-25, 1996.
- [5] Proakis, J. G., *Digital Communications*, 4th edition, McGraw-Hill, Boston, 2000.
- [6] Lehmann, E. L., *Theory of Point Estimation*, John Wiley & Sons, New York, 1983.
- [7] Ingle, R., “SAW RFID Temperature/Range System Test Results and Analysis,” NASA JSC Graduate Co-op Program Final Report, August 2008.

Appendix A

Notes on Network Analyzer Detection and Estimation for SAW RFID Applications

A.1. Introduction

Assume that a system has real-valued impulse response $h(t)$ with corresponding frequency response $H(f)$ that satisfy

$$h(t) = 0, \quad \forall t < 0 \text{ and } t > \frac{N}{B},$$

and

$$H(f) \approx 0 \quad \text{for all } |f| < f_0 \text{ and } |f| \geq f_0 + B,$$

where $B \ll f_0$ is the system bandwidth and N is the desired number of frequency samples (bins). These conditions imply that $h(t)$ can be recovered with little or no error from the frequency-domain samples $H\left(f_0 + \frac{B}{N}k\right)$ for $k = 0, 1, 2, \dots, (N-1)$. Furthermore if we have the system input-output relationship

$$\cos\left(2\pi\left[f_0 + \frac{B}{N}k\right]t + \varphi\right) \rightarrow \boxed{h(t)} \rightarrow A_k \cos\left(2\pi\left[f_0 + \frac{B}{N}k\right](t - \tau) + \varphi - \Phi_k\right)$$

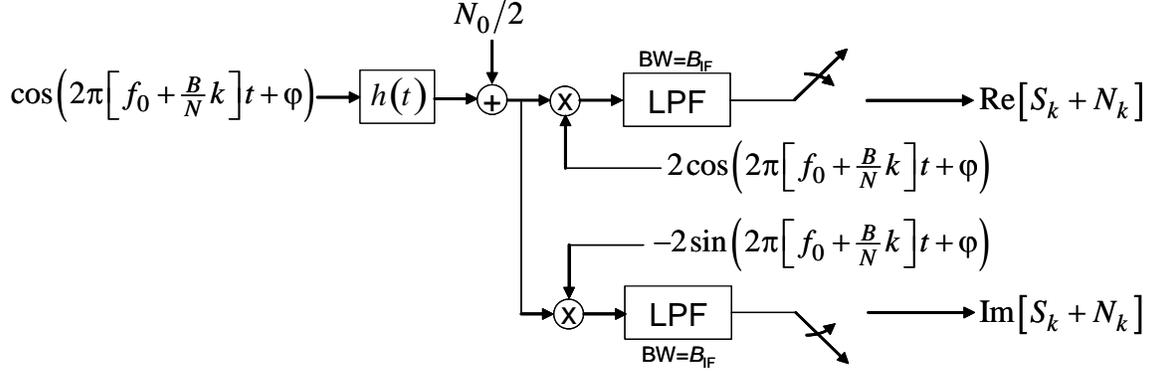
for all $k = 0, 1, 2, \dots, (N-1)$, then

$$H\left(f_0 + \frac{B}{N}k\right) = A_k e^{-i(2\pi f_0 \tau + 2\pi \frac{B}{N} k \tau + \Phi_k)}, \quad k = 0, 1, 2, \dots, (N-1),$$

and

$$\begin{aligned} h(t) &= \frac{B}{N} \sum_{k=0}^{N-1} A_k e^{-i(2\pi f_0 \tau + 2\pi \frac{B}{N} k \tau + \Phi_k)} e^{i2\pi(f_0 + \frac{B}{N}k)t} \\ &\quad + \frac{B}{N} \sum_{k=0}^{N-1} A_k e^{i(2\pi f_0 \tau + 2\pi \frac{B}{N} k \tau + \Phi_k)} e^{-i2\pi(f_0 + \frac{B}{N}k)t} \\ &= 2 \operatorname{Re} \left[\frac{B}{N} \sum_{k=0}^{N-1} A_k e^{-i(2\pi f_0 \tau + 2\pi \frac{B}{N} k \tau + \Phi_k)} e^{i2\pi(f_0 + \frac{B}{N}k)t} \right] \\ &= 2 \frac{B}{N} \sum_{k=0}^{N-1} A_k \cos\left(2\pi\left[f_0 + \frac{B}{N}k\right](t - \tau) - \Phi_k\right). \end{aligned}$$

Also, using a network analyzer to sample the system frequency response as we are doing, the output from the system can be represented in the form



for $k = 0, 1, 2, \dots, (N-1)$, where $S_k = H\left(f_0 + \frac{B}{N}k\right)$ and the N_k are i.i.d. $CN(0, 2B_{IF}N_0)$.

A.2. Signal-to-Noise Ratio

Assuming we are trying to detect the true unit impulse response $h(t)$ in the time domain in the presence of additive white Gaussian noise (AWGN) with power spectral density $N_0/2$, the signal energy is given by

$$\mathcal{E}_s = \int_0^\infty |h(t)|^2 dt = \frac{2B}{N} \sum_{k=0}^{N-1} A_k^2,$$

and the optimal detector performance for this signal is determined by the signal-to-noise ratio (SNR)

$$\text{SNR}_1 = \frac{\mathcal{E}_s}{N_0} = \frac{2B}{N \cdot N_0} \sum_{k=0}^{N-1} A_k^2.$$

If we normalize this assuming that the system is excited with a band-limited impulse having energy \mathcal{E} rather than a perfect unit impulse function, which has infinite energy, then the signal energy is given by

$$\mathcal{E}'_s = \frac{\mathcal{E}}{2B} \int_0^\infty |h(t)|^2 dt = \frac{\mathcal{E}}{N} \sum_{k=0}^{N-1} A_k^2,$$

and the normalized version of SNR_1 becomes

$$\text{SNR}'_1 = \frac{\mathcal{E}'_s}{N_0} = \frac{\mathcal{E}}{N \cdot N_0} \sum_{k=0}^{N-1} A_k^2.$$

Similarly, assuming the analyzer output Fourier coefficients S_k are known, the performance of the optimal detector based on the frequency-domain sampled signal vector $\mathbf{R} = \mathbf{S} + \mathbf{N}$, where

$$\mathbf{R} = (R_0, R_1, \dots, R_{N-1}),$$

$$\mathbf{S} = (S_0, S_1, \dots, S_{N-1}),$$

$$\mathbf{N} = (N_0, N_1, \dots, N_{N-1}),$$

is determined by the alternative SNR

$$\text{SNR}_2 = \frac{\|\mathbf{S}\|^2}{2B_{\text{IF}}N_0} = \frac{1}{2B_{\text{IF}}N_0} \sum_{k=0}^{N-1} A_k^2.$$

If we again normalize this assuming that the system is excited with time-limited sinusoids having total energy \mathcal{E} rather than pure sinusoids having infinite energy, then the normalized version of SNR_2 becomes

$$\text{SNR}'_2 = \frac{\mathcal{E}\|\mathbf{S}\|^2}{N \cdot N_0} = \frac{\mathcal{E}}{N \cdot N_0} \sum_{k=0}^{N-1} A_k^2$$

which is identical to SNR'_1 .

Finally, the performance of a simple energy detector based on the signal vector $\mathbf{R} = \mathbf{S} + \mathbf{N}$ is loosely⁴ related to the so-called *deflection ratio*, which is given by

$$\text{SNR}_3 = \frac{\|\mathbf{S}\|^4}{4N(B_{\text{IF}}N_0)^2} = \left[\frac{\sum_{k=0}^{N-1} A_k^2}{2\sqrt{N}B_{\text{IF}}N_0} \right]^2,$$

which can again be normalized to give

$$\text{SNR}'_3 = \left[\frac{\mathcal{E} \sum_{k=0}^{N-1} A_k^2}{\sqrt{N}N \cdot N_0} \right]^2.$$

A.3. Simple Estimation of individual RFID Tag Delays

Suppose that we measure the frequency response of a SAW RFID tag and we wish to estimate the individual pulse delays in the tag response directly from the frequency response. In this case the frequency response of interest is assumed to take the form

$$H_k = H\left(f_0 + \frac{B}{N}k\right) = \sum_{j=0}^{P-1} A_j e^{-i(2\pi f_0 \tau_j + 2\pi \frac{B}{N} k \tau_j + \Phi_j)}, \quad k = 0, 1, 2, \dots, (N-1),$$

⁴ For a complete discussion of the problem of SNR and error probability for energy detectors, see “On Generalized Signal-to-Noise Ratios in Quadratic Detection,” R. J. Barton and H. V. Poor, *Mathematics of Control, Signals, and Systems*, 1992. Generally, if N is large and the Central Limit Theorem holds, the deflection ratio corresponds to the usual SNR for this problem and is closely related to probability of error.

where τ_j , $j = 0, 1, \dots, P-1$, represent the *unknown* delays of the individual reflectors on the tag, and $A_j e^{-i\Phi_j}$, $j = 0, 1, \dots, P-1$, represent the *known* reflection coefficients of the reflectors. Note that with this approach, we are implicitly assuming that the reflection coefficients are independent of frequency across the entire bandwidth of the device. Changing this assumption to make the algorithm more accurate is straightforward but complicates the notation and is often neglected in practice.

To estimate the sequence of delays in a very straightforward manner, we start by reconstructing a baseband, analytic version of the impulse response given by

$$\begin{aligned}
\tilde{h}(t) &= \frac{2B}{N} \sum_{k=0}^{N-1} H_k e^{i2\pi \frac{B}{N} kt} \\
&= \frac{2B}{N} \sum_{k=0}^{N-1} \left[\sum_{j=0}^{P-1} A_j e^{-i(2\pi f_0 \tau_j + 2\pi \frac{B}{N} k \tau_j + \Phi_j)} \right] e^{i2\pi \frac{B}{N} kt} \\
&= \frac{2B}{N} \sum_{j=0}^{P-1} A_j e^{-i(2\pi f_0 \tau_j + \Phi_j)} \left[\sum_{k=0}^{N-1} e^{-i2\pi \frac{B}{N} k(\tau_j - t)} \right] \\
&= \frac{2B}{N} \sum_{j=0}^{P-1} A_j e^{-i(2\pi f_0 \tau_j + \Phi_j)} \left[\frac{1 - e^{-i2\pi B(\tau_j - t)}}{1 - e^{-i2\pi \frac{B}{N}(\tau_j - t)}} \right] \\
&= \frac{2B}{N} e^{i\pi \frac{N-1}{N} B t} \sum_{j=0}^{P-1} A_j e^{-i(2\pi f_0 (1 + \frac{N-1}{2N} B/f_0) \tau_j + \Phi_j)} \left[\frac{\sin(\pi B(t - \tau_j))}{\sin(\pi \frac{B}{N}(t - \tau_j))} \right].
\end{aligned}$$

It follows that the relationship between $h(t)$ and $\tilde{h}(t)$ is given by

$$h(t) = \text{Re} \left[\tilde{h}(t) e^{i2\pi f_0 t} \right] = x(t) \cos(2\pi f_0 t) - y(t) \sin(2\pi f_0 t),$$

where $x(t) = \text{Re}[\tilde{h}(t)]$ and $y(t) = \text{Im}[\tilde{h}(t)]$.

Note that if the delays satisfy the condition $|\tau_i - \tau_j| > 1/B$ for all $i \neq j$, then the relative maxima of the function $|\tilde{h}(t)|$ occur at the points $t = \tau_j$, $j = 0, 1, \dots, P-1$, and

$$\begin{aligned}
\tilde{h}(\tau_j) &= 2BA_j e^{-i(2\pi f_0 \tau_j + \Phi_j)} \\
&\quad + \frac{2B}{N} e^{i\pi \frac{N-1}{N} B \tau_j} \sum_{\substack{l=0 \\ l \neq j}}^{P-1} A_l e^{-i(2\pi f_0 (1 + \frac{N-1}{2N} B/f_0) \tau_l + \Phi_l)} \left[\frac{\sin(\pi B(\tau_j - \tau_l))}{\sin(\pi \frac{B}{N}(\tau_j - \tau_l))} \right] \\
&\approx 2BA_j e^{-i(2\pi f_0 \tau_j + \Phi_j)}.
\end{aligned}$$

Furthermore, if $t = \tau_j + \Delta$ where $|\Delta| < 1/B$, then

$$\tilde{h}(t) \approx \frac{2B}{N} \left[\frac{\sin(\pi B \Delta)}{\sin\left(\pi \frac{B}{N} \Delta\right)} \right] A_j e^{-i(2\pi f_0 t + \Phi_j)} e^{i2\pi f_0 \left(1 + \frac{N-1}{2N} B/f_0\right) \Delta}$$

It follows that the delays τ_j , $j=0, 1, \dots, P-1$, can be estimated by finding all of the points $t_j = \tau_j + \Delta_j$ where the observed (i.e., noisy) version of the function $|\tilde{h}(t)|$ has a relative maximum above some detection threshold and then solving for Δ_j using the relationship

$$\tilde{h}(t_j) e^{i(2\pi f_0 t_j + \Phi_j)} \approx \frac{B}{N} \left[\frac{\sin(\pi B \Delta)}{\sin\left(\pi \frac{B}{N} \Delta\right)} \right] A_j e^{i2\pi f_0 \left(1 + \frac{N-1}{2N} B/f_0\right) \Delta_j}.$$

Hence, the final estimate for each τ_j , $j=0, 1, \dots, P-1$, is given by $\tau_j = t_j - \Delta_j$, where

$$\Delta_j = \frac{\left[\arg \left[\tilde{h}(t_j) \right] + 2\pi f_0 t_j + \Phi_j \right]_{2\pi}}{2\pi f_0 \left(1 + \frac{N-1}{2N} B/f_0\right)},$$

$\arg \left[\tilde{h}(t_j) \right]$ represents the phase of $\tilde{h}(t_j)$, and the expression $[x]_{2\pi}$ represents the number x reduced modulo 2π .

In order to evaluate $\tilde{h}(t)$ at enough values of t to give high-resolution estimates of the significant relative maxima, we use the inverse FFT. We start with the vector $\mathbf{H} = (H_0, H_1, \dots, H_{N-1})$ and pad with zeroes to a length of $M > (2f_0 N)/B$, where M is a power of 2. Let $\tilde{\mathbf{H}} = (\tilde{H}_0, \tilde{H}_1, \dots, \tilde{H}_{M-1})$, where

$$\tilde{H}_k = \begin{cases} H_k, & k = 0, 1, \dots, N-1, \\ 0, & k = N, N+1, \dots, M-1. \end{cases}$$

Then, for $t = (N/MB)m$, $m = 0, 1, \dots, M-1$, we get

$$\tilde{h}(t) = \tilde{h}\left(\frac{N}{MB}m\right) = \frac{2B}{N} \sum_{k=0}^{N-1} H_k e^{i\frac{2\pi}{M}km} = \frac{2MB}{N} \left(\frac{1}{M} \sum_{k=0}^{M-1} \tilde{H}_k e^{i\frac{2\pi}{M}km} \right).$$

Hence, the vector $\tilde{\mathbf{h}} = (\tilde{h}_0, \tilde{h}_1, \dots, \tilde{h}_{M-1})$, where $\tilde{h}_m = \tilde{h}\left(\frac{N}{MB}m\right)$ for $m = 0, 1, \dots, M-1$, is a scaled version of the inverse FFT of the vector $\tilde{\mathbf{H}} = (\tilde{H}_0, \tilde{H}_1, \dots, \tilde{H}_{M-1})$.

A.4. Estimation of Temperature From Tag Delays

One approach to estimating the temperature of a SAW RFID tag, although not a very good one, is to start by measuring the individual pulse delays in the tag response for some subset of tag

pulses. The temperature of a SAW RFID tag can be estimated by measuring the tag delays if the *temperature coefficient of delay* (TCD) for the tag is known relative to a fixed reference temperature and the tag delays are also known at the reference temperature. The TCD of a tag is a function that represents the relative change in delay for any fixed reflector on the tag, with respect to the delay at the reference temperature, as a function of the change in temperature from the reference value. It can be represented equivalently as the relative change in the acoustic wave propagation velocity along the tag as a function of change in temperature, but the characterization we give here seems to be the one used most commonly in the literature.

If we let $\tau(T)$ represent the delay associated with a fixed reflector at an arbitrary temperature T , and we let $T = T_0$ represent the reference temperature for the device, then the TCD can be approximated for most purposes over fairly large temperature ranges as

$$\frac{\tau(T) - \tau(T_0)}{\tau(T_0)} \approx A(T - T_0) + B(T - T_0)^2,$$

where A and B are constants that must be determined experimentally for the particular class of devices in question at the reference temperature. Letting $\tau_0 = \tau(T_0)$, $\Delta\tau = \tau(T) - \tau(T_0)$, and $\Delta T = T - T_0$, this becomes

$$f(\Delta T) = \frac{\Delta\tau}{\tau_0} \approx A(\Delta T) + B(\Delta T)^2.$$

It follows that for any known set of reference delays $\tau_{i,0} = \tau_i(T_0)$, $i = 0, 1, \dots, P-1$, we must have

$$\frac{\Delta\tau_i}{\tau_{i,0}} = f(\Delta T) \approx A(\Delta T) + B(\Delta T)^2, \quad \forall i = 0, 1, \dots, P-1.$$

Hence, to estimate the current temperature on a tag, one approach is to first use the algorithm described in the previous section to estimate the current value $\tau_i = \tau_i(T)$ for each delay, compute the corresponding estimated change in delay $\Delta\tau_i = \tau_i - \tau_{i,0}$, and then solve the minimization problem

$$\min_{\Delta T} \sum_{i=0}^{P-1} \left[\frac{\Delta\tau_i}{\tau_{i,0}} - A(\Delta T) - B(\Delta T)^2 \right]^2$$

to find the best least-squares estimate of ΔT .

It is easy to see that there are two possible solutions to this problem, given by

$$\Delta T^* = -\frac{A}{2B} \left(1 \pm \sqrt{1 + \frac{4B}{A^2} \Delta\tau} \right),$$

where

$$\overline{\Delta\tau} = \frac{1}{P} \sum_{i=0}^{P-1} \frac{\Delta\tau_i}{\tau_{i,0}}$$

is the average relative change in delay. Evaluating the mean squared error (MSE) for these two solutions gives

$$\text{MSE}^* = \sum_{i=0}^{P-1} \left[\frac{\Delta\tau_i}{\tau_{i,0}} - A(\Delta T^*) - B(\Delta T^*)^2 \right]^2 = \sum_{i=0}^{P-1} \left[\frac{\Delta\tau_i}{\tau_{i,0}} - \overline{\Delta\tau} \right]^2 = \sigma_{\Delta\tau}^2$$

in both cases. Hopefully, only one of these two solutions makes sense physically.