Exploiting Phase Information in Pulse Doppler Radars: High Relative Resolution For The Price Of Low Absolute Resolution

Abstract

The uncertainty principle states that there is an inverse relation between bandwidth and spatial resolution. High bandwidth radars require expensive electronics and greater power, making them prohibitive for mote scale systems. Yet many applications need fine-grain resolution. In this paper, we address these contradictory requirements in the context of low power Pulse Doppler Radars, where the uncertainty is in terms of ambiguity in their absolute range information. We show that the phase of the radar returns provides non-traditional resolution information, on the scale of the wavelength, and the wavelength can be affordably made an order of magnitude finer than the traditional resolution of these systems. Paradoxically this phase information provides fine scale range information while still exhibiting ambiguity in coarse scale range information. Nevertheless, we show how to process the information to obtain much finer scale relative motion profiles than would be suggested by the uncertainty principle.

If the number of targets is small this processing can be accomplished with very simple —i.e., mote scale— algorithms. Specifically, a sequence of phase measurements can be unwrapped to directly yield an estimation of the relative motion trajectory. In the presence of noise, phase unwrapping errors create a multi-modal error distribution that must be considered in the design of algorithms and applications that exploit this result. We show that while more sophisticated algorithms can greatly reduce the frequency of phase unwrapping errors, even the simplest of phase unwrapping algorithms yields results that are easily exploited for many classification and tracking tasks that could not be performed with lower resolution sensors. As examples, we show a method for differentiating humans walking through a scene from bushes blowing in the wind by estimating the total displacement of the target, and a method for accurate network-based tracking of moving objects.

Keywords

Wireless sensor networks, pulsed doppler radar, range resolution, robust detection, fine-grain tracking

1 Context

There is significant interest in low power Wireless Sensor Networks (WSNs) [1]. An important economic and operational need is the need to operate the WSN nodes using affordable battery technology. Simple analysis of this constraint reveals a severe tradeoff between system life and average power consumption, which is illustrated in Figure 1. Although in situ energy harvesting may in time substantially reduce the severity of this tradeoff (see [2] & [3]), it still seems likely that there will continue to be strong interest in WSNs where the nodes consume on the order of 10 mW. The class of WSN nodes called motes [1], operate approximately in this power regime, and are the focus of this paper.

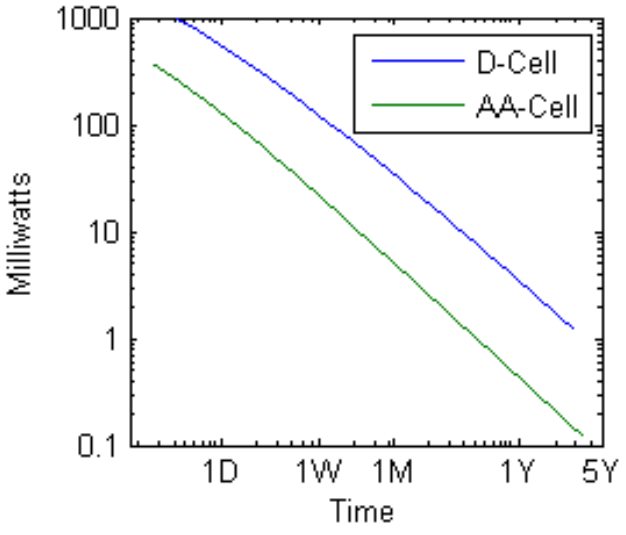

Figure 1. System life for various sized batteries.

A key obstacle to maturation of mote-scale WSNs is that they continue to struggle to overcome the limitations imposed by relatively low capability sensors. Typical mote-scale sensors include:

- Point sensors such as temperature, humidity, pressure, and light levels.
- Passive spatial and motion sensors, such as Pyroelectric Infrared (PIR) motion sensors, acoustic sensors, and still shot photo sensors.
- Active sensors, notably active ultrasound sensors and Doppler Radars.

But with few exceptions these choices offer relatively limited information content. The limitations of the information content of the available mote scale sensors has impeded the preciseness of many natural applications, such as tracking individuals moving through a WSN. In detection oriented applications, such as monitoring for perimeter protection or border patrol, the sparsity of information provided by mote-scale sensors significantly exacerbates the false alarm problem. It is anticipated that there would be many benefits to using sensors with richer information content in mote-scale WSNs. Entirely new classes of applications based on better classification might be invented. Unfortunately it is difficult to provide information rich mote-scale sensors. The key difficulty is usually power but may also include cost, and, on rare occasion, size.

This paper focuses on this tension as it applies to Pulsed Doppler Radars (PDRs). There is a secondary tension between the potential of higher resolution and the cost of electronics. But the core tension is between the information potential of higher resolution and the power required to attain higher resolution. This tension is fundamentally related to uncertainty principle which implies that greater resolution requires greater bandwidth, which in tern requires higher power (and more expensive) electronics.

Here we refer to the uncertainty principle as used in sampling theory (rather than the mathematically related uncertainty principle from quantum mechanics). There are many different formulations of the uncertainty principle; the version that seems most relevant, stated in [4], is that is that for any finite energy waveform, for a particular measure the variance of the waveform in time, denoted here by Δsig , and for a similarly defined measure of the variance of the frequency, Δf,

$$\Delta sig \geq \frac{1}{2\Delta f}$$

The temporal variance of the waveform directly affects the accuracy with which the time shift of the waveform may be robustly estimated. In the absence of noise it is possible to estimate the time offset of a signal to arbitrary accuracy, but noise creates a probabilistic limit on the accuracy with which the time shift of a signal may be estimated. If we define Δt as the accuracy of the estimate of the time offset of a waveform, then

$$\Delta t \approx k.\Delta sig$$

where the constant k depends on the Signal to Noise Ratio (SNR) and in a deterministic context on the acceptable error rate for the estimate of the temporal offset or in a probabilistic setting on the definition of Δt.

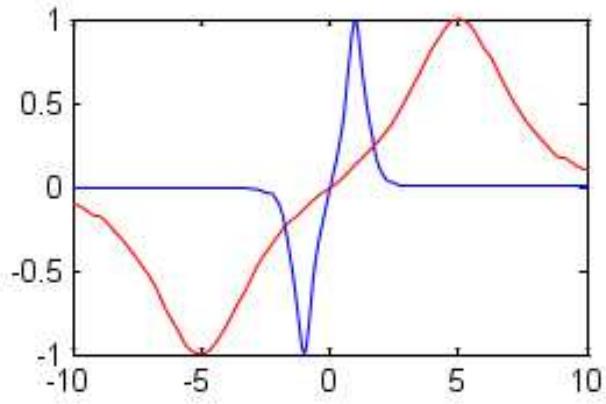

Figure 2. Two identical signals except one is scaled 5 times the other in time.

Intuition on this point may be developed by considering the two waveforms in Figure 2, one with 5 times the variance of the other. Figure 3 shows both of these signals corrupted by noise and slightly offset from the original waveform. Notice how much more the noise interferes with the ability to "see" the offset in the signal with the larger temporal variance.

Because a radar at a fundamental level estimates range by estimating differences in time-of-flight, the accuracy with which it can estimate range is directly related to the accuracy with which it can estimate time offsets of the waveform, which is directly related to the bandwidth of the waveform. If we denote the speed of light by c and the minimum achievable range accuracy by Δr, then

$$\Delta r \approx \frac{k.c}{2\Delta f}$$

for some constant k which depends on the problem formulation. It is possible to construct alternative formulations of the uncertainty principle. The history of the "classical" formulation is presented in [5]. In addition, there are quite a few more recent generalizations presented in [6]. But the conceptual point remains remarkably robust across a wide range of problem formulations: Resolution is proportional to bandwidth. And the only way to "cheat" this bothersome result is to trade SNR for resolution, which significantly increases power and/or reduces robustness.

Because greater bandwidth requires higher speed electronics, which consumes more power, there is significant incentive to do more with less bandwidth. Although the laws of physics cannot be broken, their operational implications

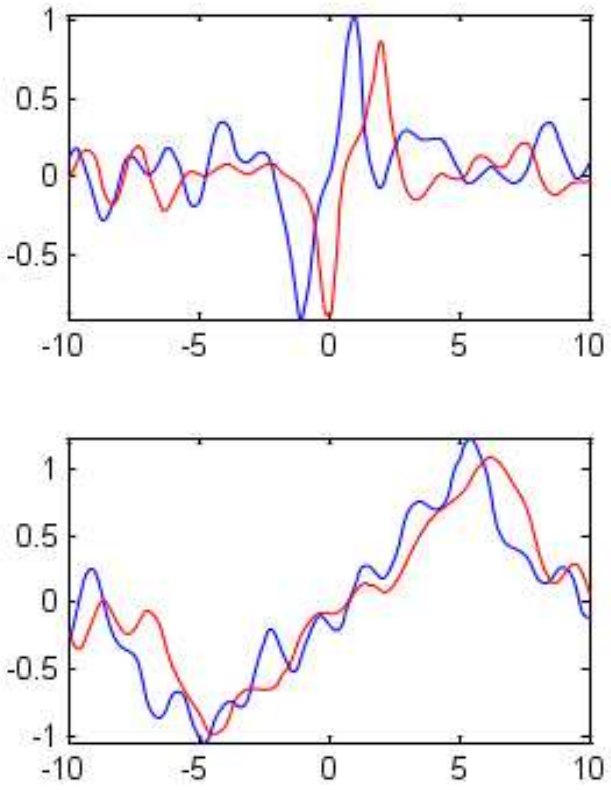

Figure 3. The same waveforms in Figure 2 but with noise.

can sometimes be circumvented. For example, [7] presents a radio interferometer that achieves accurate ranging on the order of the wavelength of the carrier even when the signal is about 4 orders of magnitude narrower. This result is achieved by exploiting the fact that, even though for normal radio operation the carrier only needs to be coherent slightly longer than the symbol duration, for extant radios the carrier is coherent for thousands of times longer than the symbol rate. This excessive coherence of the carrier allows for the exploitation of phase differences, which yield relative range information on the order of fractions of a wavelength of the carrier.

This paper presents a different method of achieving relative resolution from a coherent radar that is on the order of wavelength, even though the bandwidth is one or two orders of magnitude less than the frequency of the wavelength. We achieve this not by sacrificing[AA1] SNR, but by exploiting the relatively long time coherence of the radar, which allows us to compare the phase of returns over relatively long time intervals, e.g., one second, in order to estimate relative changes in range on the order of a wavelength. [AA1]Did you mean to say increasing SNR?

2 Pulsed Doppler Radars

In this section, we briefly overview various types of Doppler Radars, and describe elements of a PDR that is suitable for fine-grain relative resolution.

2.1 Continuous versus Pulsed Wave

The difference between traditional Doppler Radars and a Pulse Doppler Radar (PDR) is illustrated by considering a Continuous Wave (CW) radar. CW radar continuously transmits a known stable frequency while listening to echoes from the environment. It then removes the portion of the return that matches the transmitting frequency, typically by modulation of the return against the transmitted signal and applying a DC-rejection filter. The residual is the portions of the return that have some frequency shift caused by motion of the target.

From the perspective of mote scale WSNs, there are several problems with CW radars: the two that seem most important are that there are no range gates (which implies that returns from certain distances cannot be controllably rejected) and that the transmitter and receiver are on all the time. These problems can be alleviated with the PDR.

Broadly speaking, there are two types of PDRs. The first transmits pulses and computes the frequency shift between the return and the transmitted signal. This type of PDR is like using a CW radar that is turned on and off in order to create pulses. The key problem with this style of PDR is that narrow pulses yield limited frequency resolution.

The second type of PDR is more popular: it computes the change in the return between pulses in order to estimate the complement of the return which is associated with moving targets. This type of PDR requires that the returns be coherent between pulses. That is, in exactly identical environment needs to produce the same result from one pulse to the next. Coherence also implies that the signal to noise ratios can be improved by integrating over successive returns; this is useful when the signal is buried in the noise (SNR < 0, which is often the case when dealing with soft targets or small targets, such as humans, especially when the radars are deployed close to the ground) and noise in not coherent (which is typically the case). Nevertheless, real electronic systems cannot maintain perfect coherences indefinitely and the limit on which coherences can be maintained establishes a lower bound on the Doppler frequency which can be estimated. We have applied these techniques, and describe next hardware which we have developed that achieves coherences over 2 or 3 seconds.

2.2 Prototype PDR Mechanism

Although the class of PDR is slightly broader than is presented here, examining a prototypical implementation of the second type helps clarify the use of phase information. Consider a radar that transmits a series of short pulses of the type shown in Figure 4. The correlation of this pulse with itself is

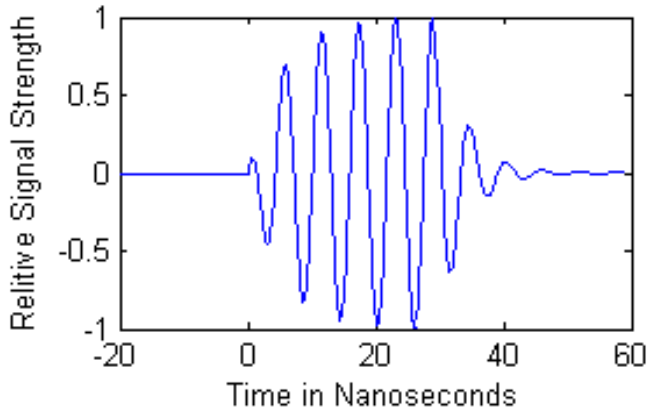

Figure 4. A prototypical pulse for a PDR.

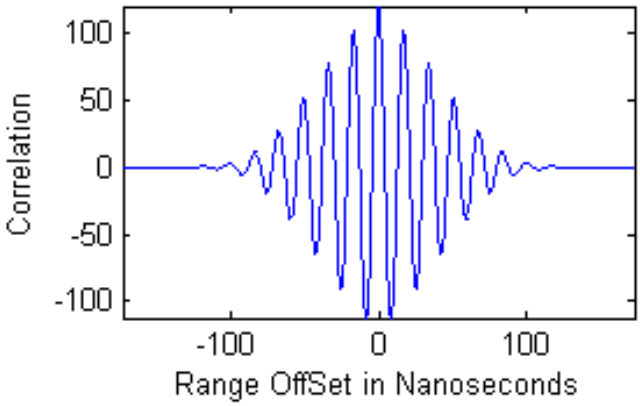

Figure 5. The autocorrelation of the pulse in Figure 4.

shown in Figure 5. The key feature of the system response of this radar shown in Figure 5 is that the response is periodic as a function of differences in flight-path length with period equal to the wavelength of the carrier frequency shown in Figure 4.

(We note that correlating the return with a copy of itself is not the only way to estimate the return from a PDR. For example, a high-end system could utilize a non-linear Impulse Response (IR) inversion technique [8], but at the time of writing this would require 2 or 3 orders of magnitude more computational power than is typically available on a mote. Other systems may employ methods for truncating correlation response in order to produce sharper range gates. This periodic system response feature is nearly universal for these PDRs as well.)

Now, consider a refinement of the PDR that lets it correlate the return with two different reference signals 90° out of phase with respect to each other, thus producing two outputs, one called the *in-phase* signal and the other called the *quadrature* signal (often denoted I & Q). The resulting outputs when using the signal shown in 4 are shown in Figure 6. These two outputs are commonly combined to form one complex sample: In cartesian terms, the I value is the real part and the Q value is the imaginary part; alternatively one may represent the sample in terms of its *amplitude* (aka absolute value/modulus) and *phase* (aka argument/angle).

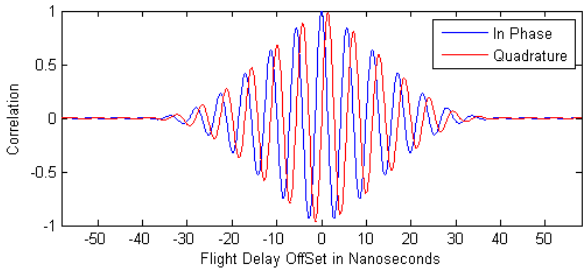

Figure 6. The output response for a "complex output" PDR.

It is common to call this type of radar *coherent*, but within the context of this paper, the overloading of the term causes confusion so we will adopt the descriptive, but non-standard, terminology *complex-output radar*.

Of course, because of the correspondence between flight delay and the range to the target, the complex output of the radar is a deterministic function of its range to the target; this is shown in Figure 7.

3 Phase Information

In this section, we first motivate why amplitude information from the complex output PDR described above is insufficient for fine-grain range information. We then explain why phase information does instead suffice, at least in a relative if not absolute sense of ranging, and describe a basic algorithm for "unwrapping the phase to reconstruct the relative motion of the target. We lastly discuss how well the basic algorithm tolerates errors due to noisy estimates of phases, and how to further improve the unwrapping when additional information (such as a motion model) is available.

3.1 Amplitude versus Phase of Complex Output

Notice from Figure 7 that the amplitude of the complex output provides only coarse-scale range information. Amplitude can differentiate whether a target is within the range bin and the range bin is on the order of the pulse width. In many operational scenarios, the target Radar Cross Section (RCS) is not known precisely, which further limits the usefulness of the amplitude information for estimating range.

By way of contrast, the phase of the output significantly limits the likely target locations within a range bin to only a small subset of the range bin as a whole. Continuing the running example, Figure 8 shows the phase of the output as a function of range. For an arbitrary phase measurement, indicated by the red line in the figure, the target could be at any

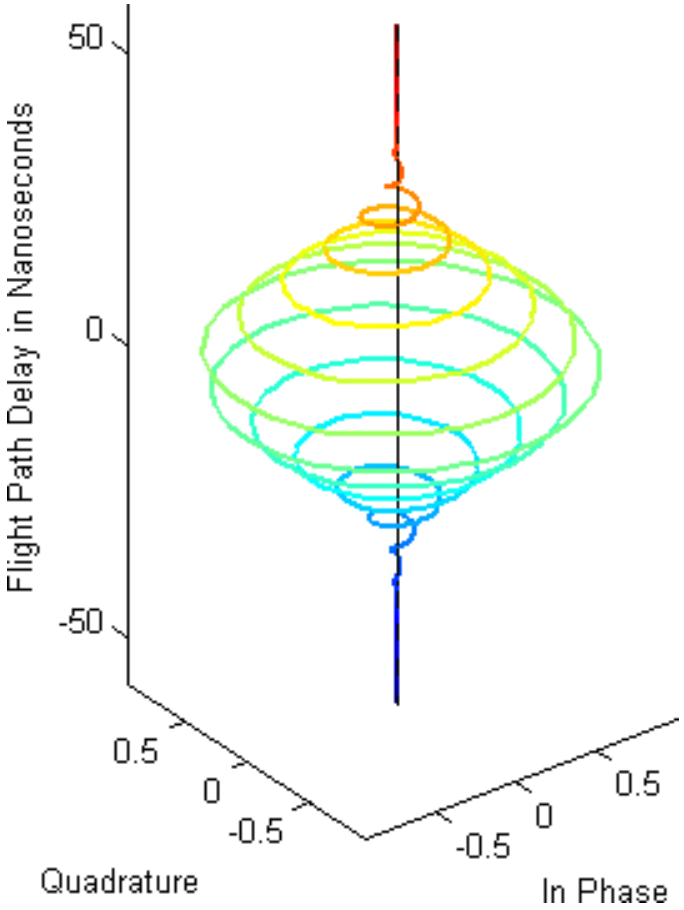

Figure 7. The display of Figure 6 as a complex real valued function.

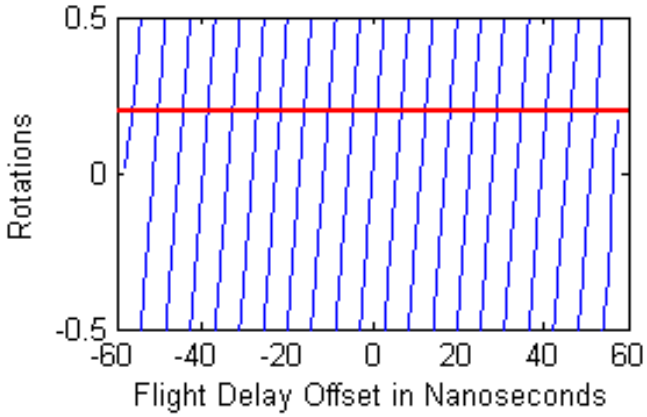

Figure 8. Phase (expressed in units of rotations) as a function of the range to the target.

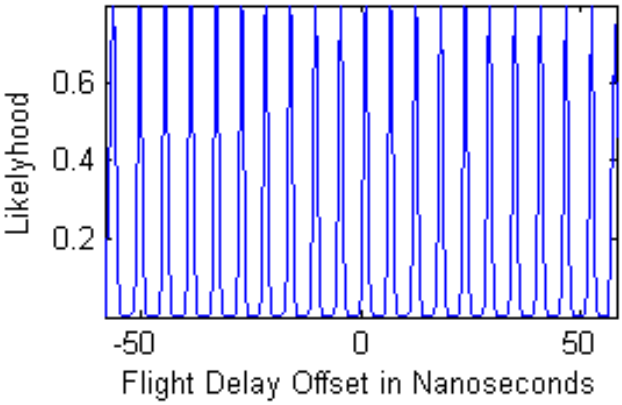

Figure 9. The likelihood function for the data in Figure 8 for an SNR of 3dB.

one of several locations, each separated by an integer number of wavelengths, but could not be in between locations. More specifically, for this phase measurement, Figure 9 shows the likelihood function for an SNR of 3 dB.

Notice that the phase measurement alone tells us with high likelihood that the target producing a return with this phase measurement is at one of roughly 20 spots within the range bin. As the SNR increases the likelihood spikes narrow and become sharper. Figure 9 uses an atypically low SNR in order to make the graph more understandable; for more typical SNRs, e.g., 6 dB or 10 dB, a single phase measurement tells us with very high probability that the target is confined to a very small percentage of the range bin, but that region is divided into 20 nearly equal sized ranges spread uniformly throughout the range bin.

By way of analogy, each phase measurement tells us that the target is somewhere on the teeth of a comb, where the width of the teeth of the comb are a function of SNR and are much thinner and have much larger gaps than for a real comb. As the target moves the "comb" moves with it. This means by observing the phase over time we can accurately reconstruct the relative motion of the target to the accuracy of the width of the teeth of the comb.

Figure 10 shows the evolution of the likelihood function as a target passes through the field of view of the complex output radar. We cannot tell which of these (roughly 20) plausible trajectories correspond to the actual target trajectory. But since all of the trajectories have nearly the same relative motion, we can know the targets relative motion to an accuracy that is a fraction of the wavelength, even when the range bin is 10 to 100 times larger.

The distortions at the edges correspond to trajectories where the target enters or leaves the range bin. For systems with multiple, partially overlapping range bins, a little bit of higher level logic can circumvent these issues.

In sum, we claim that with phase measurements and without knowledge of the amplitude (perhaps because lack of knowledge of the target's RCS renders the amplitude information unusable), it is possible to obtain *locally precise but globally ambiguous* range information.

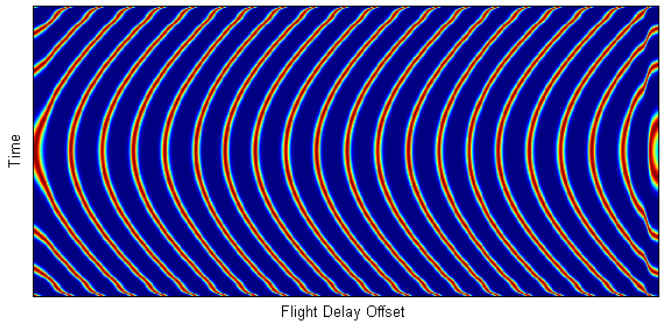

Figure 10. The evolutions of the likelihood function within a range bin as a target passes through the field of view. The red lines corresponds to maximum likelihood and the blue region to near-zero likelihoods.

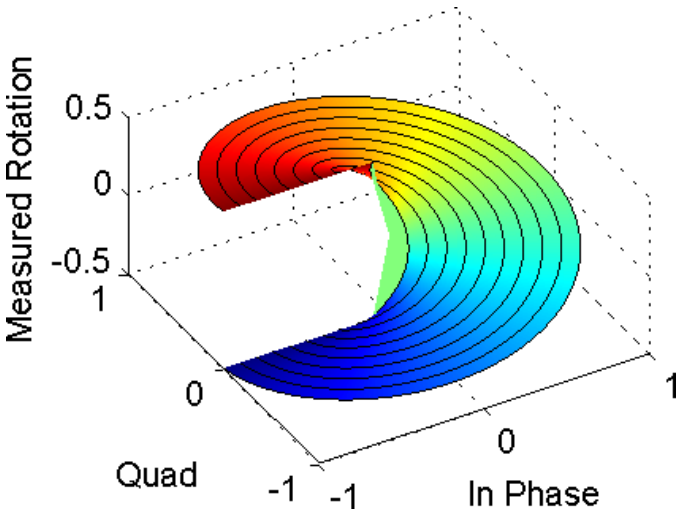

Figure 11. The standard cut in the phase function.

3.2 Phase Unwrapping

The basic computational task in phase unwrapping is to reconstruct from a sampling of the phase of the PDR output the relative trajectories, which are effortlessly seen by the human visual system in Figure 10 (or similar figures for other trajectories). At a conceptual level, it suffices for this task to measure the differences between successive phase measurements and to accumulate these pair-wise changes.

In terms of realization, a technicality has to be dealt with, as discontinuities arise in the measured phase a function of range, see Figure 8. These discontinuities arise from a discontinuity in the phase function. To better appreciate this problem, it is helpful to realize that our essential interest in the phase is to measure the amount of rotation that has occurred from an arbitrary starting point (see Figure 7). The fundamental problem is that one could traverse between any two points in the complex plane in a clockwise direction and in a counter-clockwise direction. The clockwise path will correspond to a positive total rotation and the counter-clockwise path will correspond to a negative total rotation. But these paths would differ by exactly 2π.

In complex analysis, the typical function that maps a complex value to its phase is shown in Figure 11. The cut is along the negative real axis; this choice of cut is arbitrary, but a cut is always required.

Phase unwrapping is the process of removing the discontinuities caused by the cut; this involves adding or subtracting an integer number of rotations to each phase measurement. In essence, it is the wellknown process of constructing a trajectory on the Riemann surface shown in 12.

3.2.1 Basic Algorithm

The general principle of phase unwrapping is to select at each sample point the unwrapped phase value that is closest to the expected unwrapped phase value, such that the unwrapped phase corresponds to the measured phase.

In the absence of any other information, a simple algorithm is to assume that the expected unwrapped phase is the same as the previous unwrapped phase. If we denote the measured phase at each sample by $\phi_{m,i}$ and the unwrapped phase by $\phi_{u,i}$. Then this algorithm can be stated as

$$\phi_{u,i} = \phi_{u,i-1} + mod(\phi_{m,i} - \phi_{m,i-1} - \pi, 2\pi) + \pi$$

where the sequence is seeded by $\phi_{u,0} = \phi_{m,0}$.

In the absence of noise, this algorithm exactly reconstructs the unwrapped phase if the sampling is at least at the Nyquist rate. In the presence of noise this method is quite good as long as the true rate of change of the phase between samples is small compared to half a rotation. We often use this basic algorithm even when it is not optimal, especially in settings where it is difficult to even formulate the optimal unwrapping algorithm.

3.2.2 Unwrapping Errors

In the presence of noise, the smooth likelihood contours shown in Figure 10 become distorted and occasionally they will be distorted in just the right way that adjacent lines appear to touch. In these cases, the phase unwrapping algorithm can jump to an adjacent trajectory. This is known as a phase unwrapping error. Once a phase unwrapping error occurs, the algorithm will continue along the adjacent trajectory forever, or until another phase unwrapping error occurs.

Measurements are normally perturbed by noise, so one might expect that noise would cause some accumulation of error. For instance, if instead of exploiting the phase information we were to estimate the instantaneous velocity and then integrate the velocity in order to estimate total displacement, we would expect the error on the accumulated displacement to be a random walk.

However, with phase unwrapping, small errors cancel out, unless they cause a phase unwrapping error. If we denote θ_i as the true phase and η_i as the noise induced measurement

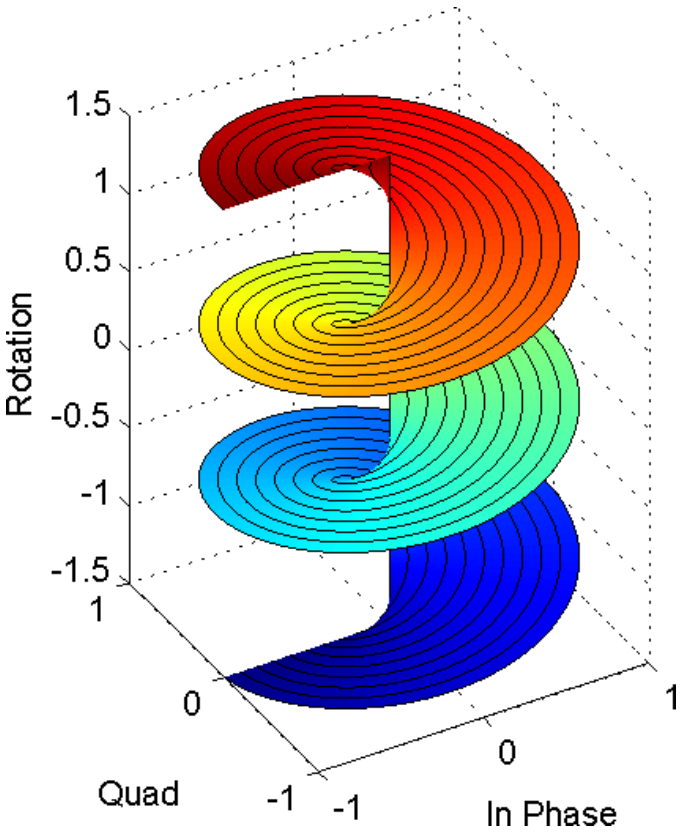

Figure 12. The Riemann surface that is the extension of the phase function shown in Figure 11.

error, then accounting for wrapping yields

$$\phi_{m,i} = mod(\phi_{m,i} + \eta_i - \pi, 2\pi) + \pi$$

It follows that if there are no phase unwrapping errors in the sequence $\{j...k\}$ then the true phase change is

$$\Delta\theta = \theta_{u,k} - \theta_{u,j}$$

and

$$\Delta\phi = \phi_{u,k} - \phi_{u,j} = \Delta\theta + \eta_k - \eta_j$$

In other words, if $\Delta\theta$ is small, the net unwrapping error is affected only by the error associated with the first point and the last point, the intermediate errors have no effect. (In case $\Delta\theta \approx 0$, it would even make sense to apply a low pass filter in order to improve the SNR until $\Delta\theta$ becomes significant.) But the performance of the algorithm deteriorates when $\Delta\theta \to \pi$. That is, when the Doppler velocity approaches the maximum detectable velocity or when the sampling rate is reduced to something near the Nyquist rate, the underlying assumptions of the basic algorithm are violated and the algorithm is more likely to jump to an adjacent trajectory.

That there is zero accumulation of error owing to intermediate points in phase unwrapping, when every phase measurement has a small amount of error, is a surprising and somewhat unintuitive result. What is happening is that routine small errors are corrected, and the cost is the occasional insertion of a larger unwrapping error.

An analogy might be drawn to financial hedging strategies that smooth out normal fluctuations, but create the risk of relatively rare catastrophic losses, except in this case the rarer large error is not catastrophic. Asymptotically, the accumulation of error occurs at a comparable rate, but rather than being spread uniformly across all samples the growth of error is bunched together in a much smaller number of unwrapping errors. This is not necessarily bad, but it is easy to build algorithm on top of phase unwrapping that assume (and depend on) a more traditional, e.g., Gaussian, error distribution.

For circular Gaussian noise in the complex output it is possible to explicitly compute the probability of a phase unwrapping error for the basic algorithm as a function of SNR and $\Delta\theta$.

3.3 Alternative Tracking Method for Phase Unwrapping

The basic algorithm implicitly exploits a random walk style motion model, which simply assumes that the next location is close to the old location. For targets with significant momentum, it would make more sense to utilize a more traditional motion model.

In this case, phase unwrapping can be formulated as a problem of phase tracking. At each iteration, the history of the unwrapped phase is applied to a motion model to predict the next unwrapped phase value. This prediction is used to select the unwrapped phase value that is closest to the predicted value and corresponds to the actual measured wrapped phase.

For instance, when applying phase tracking to track a 1-dimensional pendulum in this manner, we would use a quadratic motion model. We would perform a least squares fit of a cubic polynomial to the last (say) 30 unwrapped phase values and extrapolate one sample into the future in order to estimate the next unwrapped phase value and use this prediction to select which unwrapped phase value to use for the next point. This process is depicted in Figure 13.

Because this motion model fits the mechanics of the target, it reduces the phase unwrapping error to below the level at which we could experimentally evaluate it.

3.4 Impact of Multiple Targets on Phase Unwrapping

So far the discussion has focused on understanding the trajectories of a single target. If there is more than one target within the range bin, the complex output of the PDR will be the sum of the complex outputs that would be produced by each of the individual targets.

If it is the case that one of the targets is stronger than

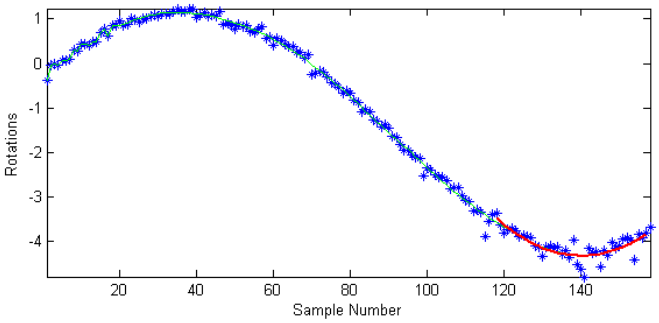

Figure 13. An example execution of a partial phase tracking algorithm. The blue points are unwrapped phase measurements, the green line is the estimated track, and the red line is the quadratic motion model that fits the last 40 points.

the rest (this is often the case with soft targets such as humans which can be regarded as a union of multiple moving objects), the interfering returns from the weaker targets add a wobble onto the trajectory of the stronger target. But because of the way that errors accumulate in the phase unwrapping algorithm, the wobble does not significantly distort the long term estimate of the trajectory of the stronger target. Of course, the presence of other moving objects in the vicinity does increase the rate of phase unwrapping errors, but if the stronger target dominates by even a few dB, this effect is not operationally significant.

For a modest number of targets with different amplitudes, if we can access a motion model, it becomes possible to track these separately. However, if these are nearly equal amplitude targets, it rapidly becomes impractical to differentiate individual targets. In effect, the unknown motion of one target acts as noise obscuring the individual motion of the other targets.

In sum, the unwrapping technique is implicitly dependent on sparsity of targets, and will not work when sensing a crowd of people. The fact that there may be a significantly stronger total return from stationary clutter than from the human tends not to be a problem because the Doppler filter removes the component of the return arising from stationary objects. In addition, the fact that there may be many other lesser moving targets in the scene tends not to be a problem because the stronger return will dominate the weaker return.

4 PDR Network Applications

This method primarily facilitates the differentiation between different types of motion rather than different types of targets. Of course in some cases a characteristic motion may be one of the more identifiable features of the target type. But more generally this technique allows us to determine whether a target is moving in a particular pattern.

Many of the mote scale PDRs are particularly well suited for monitoring human activities. The kinds of applications that this technology could enhance include detecting:

- When someone has fallen down.
- When someone is running.
- When someone is performing manual labor.
- When people are fighting.
- When someone is dancing.
- When someone is stopping at a designated location, as opposed to merely passing by without stopping.
- When someone performing an unsafe movement in a work setting.

This leads to potential applications in safety, security, activity accounting, and marketing. In addition, to this class of applications we will describe in more detail two somewhat less typical applications.

4.1 Displacement Detection

For large scale intruder detection applications of the type that might be important in civilian border patrol scenarios, facility protection scenarios, or military perimeter protection scenarios, short range sensors, especially mote scale sensors, tend to play an important role, because complex terrain often renders long range sensing incomplete. Expensive long range sensors are often too expensive to use in areas where occlusions limit the sensing range irrespective of the sensors capability. As a result it is operationally useful to deploy large scale WSN in conjunction with more traditional long range sensors.

Intruders may be readily detected with a verity of motion detectors, such as PIRs. However, a key problem with the scenario is that natural motion in the environment tends to cause false alarms. This is especially true in outdoor environment where vegetation blows in the wind. Detuning the motion detector to ignore trees and bushes blowing in the wind may render it ineffective against humans attempting to conceal their motion.

A natural solution is to analyze the type of motion in order to the motion of brush blowing in the wind from the motion of humans. Recently there has been some work on attempting to do this with PIRs that employ complicated land structures that allow them to acquire more information about the targets motion than is available through the typical PIR [**?**].

Using phase unwrapping technique presented here and a PDR, it is possible to differentiate targets that sway back and forth with negligible cumulative displacement from targets that actually move through the scene. The result is very robust and easy to implement even on low-end motes. We demonstrate this capability next through real experimental data.

4.1.1 *Experimental validation*

In this section, we describe the experimental validation of our Displacement Detection system. We present data collected in an outdoor setting using both a real human target and relevant background objects such as moving trees and bushes since the real application of our algorithm lies in robustly distinguishing background targets like trees and bushes from real targets like human intruders.

Setup. The experiments were performed using the BumbleBee Pulsed Doppler Radar [9] connected to a TelosB mote. The in-phase and quadrature radar data was sampled at 333Hz and logged to a laptop for phase unwrapping and displacement analysis.

In the first set of experiments, we set the BumbleBee radar close to a bush or a tree in moderate to high wind conditions so that the back and forth blowing of the branches and leaves could be observed. In the next set of experiments, we had a human target walk through the same scene traversing a physical distance of several meters. *Results.* Figure 14 shows the raw (amplitude) data and it can be seen that the amplitude of the radar signal goes high whenever there is motion, whether it is caused by a blowing tree or by a human target walking through the scene. Moreover, the amplitude of the radar returns is comparable for a human target and a large tree (12ft tall in this case) moving in high wind. However, Figure 15 shows the unwrapped cumulative phase for the same data and we clearly see that the unwrapped phase changes sharply when the motion traverses across a certain physical distance. To quantify this difference, we plot the frequency distribution of the total unwrapped phase change over a fixed time interval for background and human targets, which is shown in Figure 16.

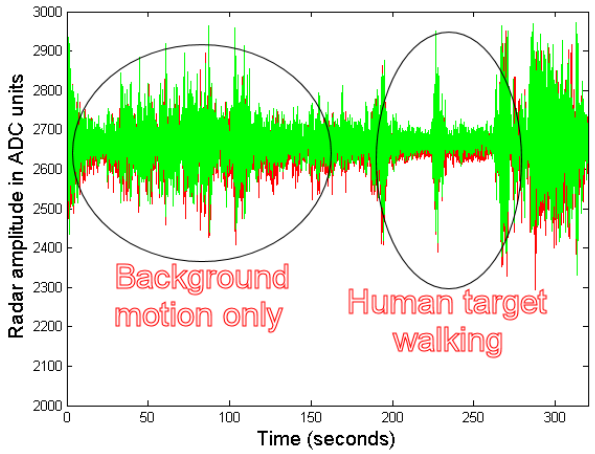

Figure 14. Raw data for tree blowing in the wind and human target walking through the scene.

We see from Figure 16 that there is a clear separation in the phase change over a certain time window (1 second in this case) between motion that is concentrated at the same

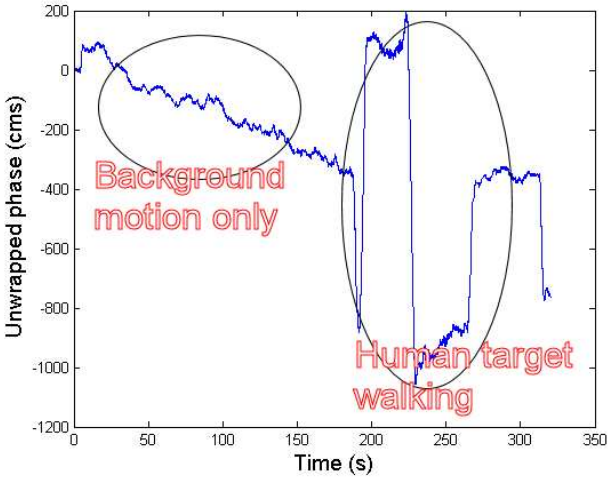

Figure 15. Unwrapped phase.

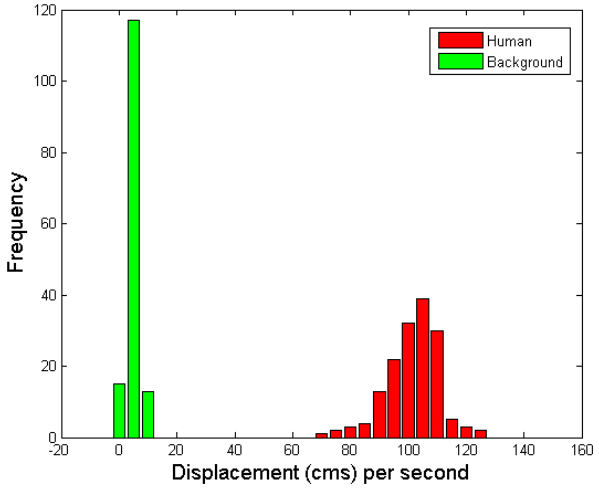

Figure 16. Histogram of cumulative displacement for background vs. human targets.

location and motion which traverses a physical distance in space. The histogram allows us to clearly select a spatio-temporal threshold (60 cms per second in this case) that allows us to robustly distinguish motion of a human target that traverses across physical space from that of background objects like trees and bushes which is localized.

5 References

[1] Jennifer Yick, Biswanath Mukherjee, and Dipak Ghosal. Wireless sensor network survey. *Comput. Netw.*, 52(12):2292–2330, 2008.

[2] Kris Lin, Jennifer Yu, Jason Hsu, Sadaf Zahedi, David Lee, Jonathan Friedman, Aman Kansal, Vijay Raghunathan, and Mani Srivastava. Heliomote: enabling long-lived sensor networks through solar energy harvesting. In *SenSys '05: Proceedings of the 3rd international con-*

ference on Embedded networked sensor systems, pages 309–309. ACM, 2005.

[3] Prabal Dutta, Jonathan Hui, Jaein Jeong, Sukun Kim, Cory Sharp, Jay Taneja, Gilman Tolle, Kamin Whitehouse, and David Culler. Trio: enabling sustainable and scalable outdoor wireless sensor network deployments. In *IPSN '06: Proceedings of the 5th international conference on Information processing in sensor networks*, pages 407–415. ACM, 2006.

[4] Ahmed I. Zayed. *Advances in Shannon's Sampling theory*. CRC Press, 1993.

[5] G. B. Folland and A. Sitaram. The uncertainty principle: A mathematical survey. *Journal of Fourier Analysis and Applications*, 3(3):207–238, May 1997.

[6] J.A. Hogan and J.D.Lakey. Embedding and uncertainty principles for generalized modulations space. *Modern Sampling Theory: Mathematics and Applications*, pages 75–108, 2001.

[7] Miklós Maróti, Péter Völgyesi, Sebestyén Dóra, Branislav Kusý, András Nádas, Ákos Lédeczi, György Balogh, and Károly Molnár. Radio interferometric geolocation. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 1–12. ACM, 2005.

[8] Kenneth J. Hintz. Snr improvements in niitek ground-penetrating radar. volume 5415, pages 399–408. SPIE, 2004.

[9] The Samraksh Company. BumbleBee Mote-Scale Pulsed Doppler Radar. `http://www.samraksh.com/products.htm` .