False Alarm Management using Phase Information for Displacement Detection

Kenneth W. Parker, Sandip Bapat and Anish Arora

Abstract

If each sensor acts independently, then dense Wireless Sensor Networks (WSNs) rapidly hit fundamental performance limits associated with False Alarm Rates (FARs). In order to be viable, WSNs must use local redundancy to manage the system-level FAR. Because a small amount of redundancy has substantial impact, this transforms the false alarm problem from a noise-centric problem into a Target Not of Interest (TNI) problem. The solution to the TNI problem is better target classification.

The most frequent, and therefore most fundamental, TNI problem in outdoor environments is vegetation that moves in the wind. Displacement Detection addresses this by differentiating targets that move back and forth from targets with non-negligible net displacement. Obtaining displacement via high resolution absolute range measurements requires wide bandwidth, which is expensive and consumes too much power. Instead, we use phase information from a far lower power and less expensive medium bandwidth Pulse Doppler Radar to obtain profiles of relative range with respect to time. We demonstrate a low complexity algorithm for using this information to perform Displacement Detection.

Although this method may not work for much longer ranges (e.g., kilometers) due to the larger number of moving objects within the scene, we demonstrate that it works well in a wide range of realistic environments at short ranges (e.g., 10's of meters). In addition, we analyze the FAR and the Probability of Detection, showing that this method reduces the FAR in most environments to about one false alarm every 3 months, even without the use of network-level redundancy, which is still necessary for dense WSNs.

Keywords

Wireless sensor networks, false alarm management, displacement detection, pulsed doppler radar, phase exploitation, interferometric radar

SenSys'10, November 3–5, 2010, Zurich, Switzerland

1 Introduction

A key thesis of this paper is that elevated False Alarm Rates (FARs) are one of the most important under-addressed problems in Wireless Sensor Networks (WSNs). We propose that this, inherently context specific, thesis holds for a large class of intrusion detection problems as well as human activity monitoring problems where depth of coverage and battery powered mote-scale sensors are required. However, because redundancy rapidly eliminates false alarms due to noise, the true cause of FAR problems in a well designed system are Targets Not of Interest (TNIs).

The main contribution of our work is to address the problem of false alarms induced by motion of background objects such as bushes and trees that are quite common in outdoor WSN deployment environments. We reject the motion of these objects, which move back and forth in the same spot, by using net target displacement as the feature for target detection.

We identify Phase Unwrapping using a Pulse Doppler Radar (PDR) as a basis for realizing Displacement Detection with low power. Phase information is used to calculate relative range profiles, which are all that is required for Displacement Detection. The alternative of using absolute range information with a range resolution of a few centimeters would require that we use a radar that has a bandwidth of a few GHz. However, state-of-the-art radars with such a wide bandwidth consume higher power than is appropriate for mote-scale devices. In contrast, relative range resolution only requires that the maximum frequency of the radar exceeds that many GHz, which is feasible today with low-power radars.

We moreover show how Displacement Detection can be implemented locally on mote-scale devices using a low-cost, low-power, Commercial off the Shelf (COTS) radar and validate its performance through rigorous experimentation in different deployment environments. Our algorithm achieves a rate of 1 false alarm every 84 days with a detection probability of 1, thereby making it especially appropriate for dense or large WSN deployments.

We note that the WSN community has successfully experimented with low-power PDRs before. Dutta et al [1] report on deploying a network of PDR motes that used a simple Neyman-Pearson binary hypothesis detector for discriminating targets from noise but not moving background objects. The realization we report here was successfully used by the winning team in a recent Commander's Challenge competition, which involved detecting intruders in a mountainous

border region [2].

The rest of the paper is organized as follows. We set up the TNIs problem for the context of people detection in outdoor environments. We describe the phase unwrapping based Displacement Detection in Section 3. Although not central to the paper, the key concepts of Pulsed Doppler Radar technology and the details behind the use of phase information may be of interest; these are discussed in Appendix A.1 and Appendix B. We discuss our experimental methodology and describe the results that validate our algorithm in Section 4. We discuss implementation considerations for low-power, mote-scale devices in Section 4.3 and conclude in Section 5.

2 Problem Formulation

2.1 A Motivating Application

The context of this paper is a large, but specific, class of applications–detecting people in relatively open spaces. These applications are central to various military, facility utilization, safety, and security domains. Militarily relevant scenarios include perimeter protection [3]. Other scenarios include customer monitoring in retail spaces, construction site safety monitoring, activity monitoring in ship yards and factory floors, auditing the use of civic spaces, and recognizing when visitors to wilderness parks wander into dangerous areas. The results of this paper are relevant to this wide range of applications, but in order to be more concrete we will elaborate on the perimeter protection scenario as a kind of canonical embodiment of this class of problems.

Imagine a perimeter where it is important for security or political reasons to detect individuals crossing the perimeter on foot. A natural question is, will a thin line of sensors along the perimeter suffice or is it necessary to deploy sensors deeper within the protected territory. The appropriate depth of coverage depends on several factors, including the accuracy of detection as well as the achievable and allowable response time. Some choices of sensing modality require distributed detection within a neighborhood of nodes, low tolerance to false alarms may lead to the use of node level redundancy, and nontrivial response times (say in the case of human responders) would imply that sensor coverage be deep enough to retain surveillance during the response time.

2.2 Low Power Wireless Sensor Networks

When providing depth of coverage in outdoor or otherwise open locations, economic and operational viability usually depends on the sensors operating for long periods on battery power. In such scenarios, dense access to a power distribution network or central power generation is usually impractical. Furthermore, as the number of sensors increases, changing the batteries becomes operationally burdensome. For example, changing the batteries on a network of 100 sensors once a month requires a small team of technicians to travel around the network, changing on average 4 batteries per day, without revealing the sensor locations to potential intruders. For a larger network, the batteries must typically last at least a year to be operationally relevant.

However, for representative batteries, this implies a severe constraint on average power consumption. Figure 1 illustrates the tradeoff between system life and power consumption. As technology matures, in situ energy harvesting may substantially reduce the severity of this tradeoff, [4] & [5], but there is, and will continue to be for many years, a strong interest in WSN nodes that consume 10s of milliwatts. Devices in this class are conventionally referred to as *mote-scale* [6].

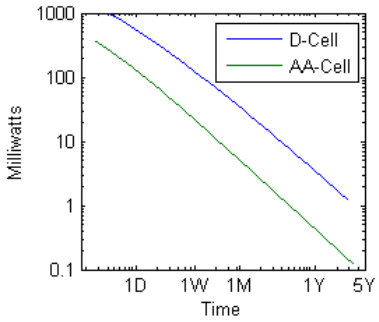

Figure 1. System life for various sized batteries.

We therefore focus on the use of mote-scale radars in the intrusion detection problem. The FAR problem is closely related to the performance limitations of mote-scale sensors: the need to use 10s of milli-watts, on average, including sensor specific computing, is the prominent cause of relevant performance limitations, and is more constraining than either cost or form factor.

2.3 False Alarm Rate Problems with Scaling

Experiments and demonstrations of WSNs using mote-scale sensors for people detection have struggled with FAR problems. The fundamental cause of this problem is that, if each node-level alarm maps directly to a system-level alarm, then acceptable node-level FARs tend to result in unacceptably high system-level FARs. That is, if the nodes act independently for the purpose of detection, then the system-level FAR is much higher than the node-level FAR, and it is easy to design a system where the node-level performance is acceptable, but the system-level performance is not.

As an arbitrary example, consider a sensor that makes a detection decision once every second. Let P_f denote the probability of false alarm and let P_d denote the probability of detection. If for an algorithm/sensor pair $P_f = 10^{-6}$ and $P_d = 0.9$, then on average a node would generate a false alarm once every 11.6 days. In most perimeter protection scenarios, this would be a good result. However, if $3 \cdot 10^4$ such nodes were deployed in a single system, then, on average, one node somewhere in the system would experience a false alarm once every 34 seconds, rendering the system unusable, at least for most scenarios.

For mote-scale systems, this problem is exacerbated by the relatively weak performance of the individual sensors. Many mote-scale sensors don't normally achieve $P_f = 10^{-6}$.

2.4 The Power of Redundancy

In a very real sense this example is a failure to utilize one of the key potential advantages of dense WSNs, namely redundancy. Consider the case where 3 sensors can sense the target and each sensor makes a local detection decision independent of all the other sensors. Then the system designer

may choose to declare a system-level detection if any one of the sensors declares a local detection, to declare a system-level detection only if all of the sensors declare a local detection, or for some in-between rule. If we denote the group-level probability of false alarms by $P_{f,g}$ and the group-level probability of detection by $P_{d,g}$, then for an m-out-of-n rule:

$$P_{f,g} = \sum_{k=m}^{n} \binom{n}{k} P_f^k \cdot (1-P_f)^{n-k}$$
$$P_{d,g} = \sum_{k=m}^{n} \binom{n}{k} P_d^k \cdot (1-P_d)^{n-k} \quad (1)$$

Equation 1 implies that the group-level decision is much more powerful than local decisions. For example consider a sensor with $P_f = 10^{-4}$ and $P_d = 0.9$. By traditional sensor design standards this is a fairly low quality detector; at one decision per second it averages a false alarm every 2.8 hours. As a stand alone sensor this would be marginal for most perimeter protection scenarios, but it is perhaps representative of many mote-scale sensors. Table 1 shows the result of various decision criteria for this fictitious sensor. The group-level FAR decreases exponentially with the amount of redundancy, so a very small amount of redundancy reduces the FAR to levels where it becomes unimportant.

Table 1. System level effect of various m-out-of-n rules, for $P_f = 10^{-6}$, $P_d = 0.9$, and $n = 3$.

m	$P_{d,g}$	$P_{f,g}$	One Group	$3 \cdot 10^4$ Groups
1	0.999	3e-4	56 min	110 ms
2	0.972	3e-8	1.1 yr	19 min
3	0.729	1e-12	32e3 yr	1.1 yr

A system wide redundancy ratio of 3 might seem excessive if the only source of redundancy were spatial redundancy. However, there are many other sources of redundancy. The most natural would seem to be the exploitation of spatio-temporal patterns. So it is not unreasonable to achieve in excess of 3 fold redundancy. For this reason great sophistication in the exploitation of redundancy is usually not required. A little bit of redundancy used somewhat efficiently, makes the classical FAR problem negligible.

2.5 The Target Not of Interest Problem

Of course, a key premise of this paper is that the FAR problem is not quite this simple. Many WSNs exploit redundancy but still struggle with FAR problems. The problem is not that the theory is flawed, but rather with the applicability of the theory. The argument for the power of redundancy implicitly assumes that false alarms are caused by events that are independent from one sensor to next. This is true for noise-induced false alarms. This source of false alarms is nearly eliminated by almost any use of redundancy. However, there is another source of false alarms that are highly correlated across local groups of sensors.

Before we define this other source of false alarms, let us consider the example problem of detecting Sport Utility Vehicle (SUV) tire noise acoustically.[1]

Tires are designed to make approximately random noise over a range of frequencies in order to maximize passenger comfort. Experimentation shows that two different features can be used to detect tire noise with some effectiveness: 1) loudness over the appropriate frequency range and 2) the shape of the acoustic spectrum, [7, 8]. Loudness, even over a specific frequency range, may seem like an inappropriately simple feature, yet it works surprisingly well most of the time. However, its lack of specificity means that a song bird, for example, could cause false alarms. And redundancy would not remove this source of false alarms because there is a real sound that is confused with the sound of a SUV. The spectral shape feature might avoid this source of false alarms, but its computational complexity is much higher.

This sort of trade-off between specificity of the detector and the computational cost is nearly ubiquitous. As a result, in severely resource constrained designs it is almost always the case that the detection class is somewhat larger than the class of operationally relevant targets. If the operational need is to find all Toyota Land Cruisers, the system might actually detect all SUVs and employ resources outside the WSN (e.g., a human) to decide what type of SUV was detected. We will define TNIs as objects which are detected under the detection model, but which are not of operational interest.

The point of this section is that proper exploitation of redundancy, even modest redundancy, exploited sub-optimally, eliminates nearly all noise-induced false alarms and that the remaining false alarms are almost all TNIs. That is, exploitation of redundancy transforms the FAR problem from a noise-centric problem into a TNI problem. This problem is under-addressed in the WSN literature.

In order to avoid running overly costly detectors, we propose a multi-stage detector, where the first stage employs only the lowest "cost" features, secondary (or tertiary) detection filters based on higher cost features are invoked infrequently. For historical reasons, these secondary detection filters are usually referred to as classifiers. We propose a low-power classifier to eliminate the most prevalent source of TNIs, in the outdoor setting, i.e., vegetation blowing in the wind, next.

3 Problem Solution

Most mote-scale intrusion detection WSNs use motion detectors that can sense the presence of motion but provide minimal information about the details of the motion. Moving brush and vegetation in an outdoor environment tend to create an intractable TNI problem for these systems. Utilizing features based on net-displacement rather than motion appears to be the best way to eliminate this source of false alarms. However, this requires a profile of target range with respect to time that is accurate to less than 5 cm.

Achieving 5 cm of absolute range resolution requires about 3 GHz of bandwidth which implies an Ultra-Wide Band (UWB) radar, and requires more power than is compatible with a mote-scale implementation. However, 5 cm of relative range resolution only requires that the maximum frequency be greater than 3 GHz, it makes no requirement on

[1] For most mass-market non-commercial vehicles, the engine noise is sufficiently muffled so that tire noise is a more detectable acoustic feature of the vehicle.

bandwidth. We use a COTS PDR [9], with about 100 MHz of bandwidth and a center frequency of 5.8 GHz in order to create a relative range profile for displacement detection.

Creating the relative range profile is as simple as unwrapping the phase of the radar's output. This results in two types of relative range errors: 1) noise errors and 2) phase unwrapping errors. We used a pendulum to estimate the accuracy of the resulting relative range profile and found that the noise error was less than 5 mm and that using the most elementary phase unwrapping algorithm, phase errors occurred, on average, less than once every 10 s. Each phase unwrapping error results in target range error of exactly $\pm\lambda/2 = 2.6\,\text{cm}$. So about 6 times a minute, the relative range profile experiences a jump of $\pm 2.6\,\text{cm}$. We were unable to determine ground truth for a swaying bush or a walking human to this kind of accuracy, but we believe that the resulting relative range profiles are very accurate.

3.1 Feature Design

3.1.1 Motion Detection Outdoors

Most mote-scale people detection systems are more humble than the previous subsection might suggest; typically these systems employ sensors such at PIRs that are optimized for motion detection and offer minimal support for features other than simple indications of motion. In some indoor environments, motion detectors work well, because almost nothing moves unless a human is present. However, in most outdoor environments, especially in more remote settings, the environment may exhibit a great deal of motion even when no humans are present.

Our experimentation strongly suggests an approximate hierarchy of TNIs that affect motion detectors in remote settings (from most common TNI to most difficult to classify):

1. Brush and vegetation blowing in the wind.
2. Small wildlife, such as birds and rodents, that come arbitrarily close to one of the sensors.
3. Large wildlife, such as deer and pigs, that wander through the system.

The last of these problems is quite difficult to solve. Designing any sensor/algorithm pair that can robustly distinguish a wild pig from a human, for example, is challenging, even without worrying about mote-scale resource constraints. Fortunately, this problem is rare enough that it may either be safely ignored or the particular situation may allow for some use of heavyweight (i.e., non mote-scale) devices.

The second problem can be addressed through the same type of group-level fusion that was discussed in Section 2.4 to reduce the random sources of false alarms. A grasshopper on the surface of one sensor may cause problems that are nearly intractable with that sensor, but the grasshopper will not be seen by any of the neighboring sensors, and the odds of several sensors all being confused in a mutually consistent way at the same time is small.

Since the first of these problems is nearly ubiquitous, it is the most fundamental of the lot. It must be solved efficiently, using only mote-scale resources.

3.1.2 Displacement Detection

The feature that most easily, and most reliably, distinguishes brush blowing in the wind from all the objects with more plausibility as operational targets, seems to be that brush blows back and forth with negligible net-displacement while all of the more challenging targets of operational interest move though the scene with non-trivial net displacement.[2] As a result, we propose a family of detectors based on measurements of net displacement.

When a bush sways in the wind, it exhibits oscillating negative and positive motion (from the sensors perspective) that exactly cancel out over time. In practice, errors in sensing the motion cause the measurements of motion to not exactly cancel out over time. This accumulation of motion measurement errors results in the perception that a swaying bush is actually engaged in a very slow random walk. This practical consideration prevents a pure displacement detector, rather the practical implementation looks for net displacement to exceed what would be expected by the random accumulation of motion measurement errors. The simplest form of this approach is simply looking for displacement to exceed some threshold within some specified time interval.

3.2 Estimating the Range Profile to High Accuracy

A key problem with implementing displacement detection is that it requires high precision profiles of the target range as a function of time. There is no particular resolution that is necessary for robust displacement detection: more accuracy results in lower FAR and higher Probability of Detection (PD). A useful approximation is that we need about 5 cm of accuracy in order to get good results, otherwise the accumulation of errors becomes large.

3.2.1 Ranging Radars and Absolute Range Resolution

Perhaps the most classical way of constructing a range profile is with a ranging radar. The process is:

1. Repeatedly measure the distance to all targets within the field of view.
2. Associate targets in one measurement, usually referred to as a scan, with targets in other measurements, possibly with the aid of a sophisticated motion model.
3. Use the sequence of associated target measurements to construct the target's range profile, from which displacement detection can be performed.

Such a system can be made to work, but has to address three somewhat serious issues:

1. The returns from the large amount of stationary clutter within the field of view must be learned and mathematically removed in order to sense the smaller number of moving targets.
2. Achieving adequate range resolution requires a lot of bandwidth, which is expensive and consumes a lot of power, more than what seems compatible with battery powered operations.

[2] We choose to define a human that moves back and forth while remaining in one spot as operationally uninteresting.

3. Radio Frequency (RF) regulatory regime somewhat strongly discriminates against this approach.

The first issue may be addressed as follows. A human sized target walking through a vegetated area may create a smaller radar return than the total return from all the clutter. As a result, the target does not stand out against the background and must be found by learning the expected return from stationary clutter and subtracting this from the instantaneous return. The residual reveals all of the smaller targets that are not normally part of the environment. The body of methods used for this process is referred to as Change Detection [10–12]. This process is more complex than the solution presented in this paper.

The second issue, i.e., achieving adequate range resolution, is more fundamental. Let $\triangle R_a$ denote the absolute range resolution and $\triangle f$ denote the bandwidth. Then an information theory result states that[3]

$$\triangle R_a \geq \frac{1}{2 \cdot \triangle f}. \quad (2)$$

Note that this result applies to the flight path length, so target range resolution is half this (or twice as good). Now, it is possible to buy a COTS UWB radar with sufficient bandwidth to achieve our results. For example the PulsON radar from Time Domain Corp., [13] achieves absolute range resolution approaching 1 cm. [4] However, it appears that achieving this resolution with a power level that is compatible with mote-scale systems is beyond the state of the art and that mote-scale COTS radars with 3 to 5 GHz of bandwidth will not be available anytime soon. For example, in late 2009 the PulsONs used from 2.5 to 6 W and cost several thousand dollars each.

This issue is not fundamentally related to radiated energy or the analog electronics. For example, the PulsON radiates less than 1 mW of power. Rather, the problem is the amount a computing that is required to process such large amounts of data. A radar requires at least 4 or 5 billion arithmetic operations per second for each 1 GHz of bandwidth, and sophisticated radars might require much more than that. The power requirements of the associated signal processing make 3 to 5 GHz radars not mote-scale even if the radar power went to zero.

The final issue relates to the regulations that govern licensing of UWB radars, because these radars emit radiation across a huge swath of the spectrum, including portions of the spectrum that are licensed for the exclusive use of organizations that paid for that right. As a result, the allowed energy densities are quite low. For example, in the United States of America (USA), UWB radars are limited to 100 nW/MHz, [14]. At 1 GHz of bandwidth this adds up to 100 W of total power. (By comparison, the Industrial Scientific and Medical (ISM) band regulations, which apply to say wireless keyboards, allow up to 1 W, which is 10,000 times as much total power. These low power levels map directly into low SNRs.

3.2.2 Phase Information and Relative Range Resolution

At first glance, a non-specialist reader might be surprised at the use of bandwidth instead of maximum frequency in the statement of Equation 2. The role of maximum frequency is evident in a related result. Let $\triangle R_r$ be defined as the relative range resolution, i.e., the resolution at which changes in range can be discerned, and let f_{max} be defined as the maximum frequency, then [5]

$$\triangle R_r \geq \frac{1}{4 \cdot f_{max}}. \quad (3)$$

For a narrow bandwidth system the absolute range resolution may be very large (i.e., very poor) compared to the relative range resolution, because the bandwidth is small compared to the maximum frequency.

Interferometers were among the first practical applications of the distinction between Equation 3 and Equation 2, [15]. Scientists realized that they could measure surface textures, and other features, with relative accuracies on the order of the wavelength of light, even though they had no way (with the technology of that era) to measure absolute range on such a fine scale. Interferometric principles were soon applied to acoustic systems ([16]), radars ([17, 18]), and lasers [19]. More recently, they have found application in mote-scale WSNs radios [20, 21].

The basic idea behind interferometric techniques is that change or differences in the phase of returned signal corresponds to changes or differences in the path length on the order of a fraction of wavelength. Most interferometric radar work has focused on difference in path length between multiple transmitters or receivers. In this paper, we use the same principle to measure small changes over time in path length resulting from the motion of the target. This requires that we measure these changes often enough so that the target has moved less than half a wavelength between measurements. But by measuring the phase often and tracking the changes in the phase over hundreds of measurements, we can construct a relative range profile with an accuracy $\lambda/8$.

This alternative is surprisingly well suited for this application, because: 1) relative range profiles are adequate for displacement detection and 2) it is relatively easy to build a medium bandwidth mote-scale radar with a center frequency in the 3 to 5 GHz range.

3.2.3 Radar Platform

Medium bandwidth mote-scale Doppler radars have been widely used in research for about 15 years. We adopted the BumbleBee radar [9] as it is a COTS device, it provides phase information (the Advantaca TWR-ISM-002 radar of [1] does not), and it interfaces with existing motes.

[3] This result holds for a specific assumption about the Signal to Noise Ratio (SNR). A precise statement that accounts for the ability to trade SNR for resolution, in a process that is called super-resolution, becomes complicated and is beyond the scope of this paper.

[4] The 3 dB envelope resolution for the PulsON is about 7.5 cm, but our experimentation shows that this radar has excellent repeatability and that if the SNR is high, sophisticated processing of the output achieves a resolution approaching 1 cm.

[5] The footnote for Equation 2 also applies to Equation 3.

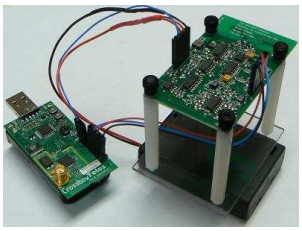

Figure 2. The BumbleBee radar.

Shown in Figure 2, the BumbleBee has a range of 10m with one range bin, which implies a range resolution of 10m. Its bandwidth is less than 100 MHz, which conforms to the FCC's ISM regulations. Its center frequency is 5.8 GHz, which implies a wavelength of 5.17 cm, and Equation 3 suggests that we should expect relative range accuracy of about 6.5 mm. And it consumes 41 mW of power.

3.3 Algorithmic Details

The process of tracking phase changes in order to construct a representation of total change from an arbitrary starting point is known as *Phase Unwrapping*. It is presented more formally in the appendix.

3.3.1 Basic Algorithm

If we denote the change in target range from an arbitrary starting point at sample times t_i by $\triangle r_i$, the measured phase by $\phi_{w,i}$, and the measurement error by η_i, then

$$\phi_{w,i} = \mod(\triangle r_i\,(2\pi/\lambda) + \eta_i + \pi, 2\pi) - \pi. \qquad (4)$$

We further define *unwrapped phase* as,

$$\phi_{u,i} = \triangle r_i\,(2\pi/\lambda) + \eta_i.$$

Phase unwrapping is roughly the process of inverting the lossy measurement process; specifically, it is the process of computing the set of $\phi_{u,i}$'s from the set of $\phi_{w,i}$'s. Notice this is subtly different from estimating the set $\triangle r_i$, because the result of unwrapping retains the measurement errors η_i but removes the effect of the modulus operation in 4. Stated differently, phase unwrapping is the process of selecting the set of $k_i \in Z$ such that

$$\phi_{u,i} = \phi_{w,i} + k_i \cdot w\pi,$$

so that $\triangle r_i = (\lambda/2\pi)(\phi_{u_i} - \eta_i)$.

The unwrapping problem is only tractable if $\triangle r(t)$ has some smoothness properties that can be exploited. Ideally, $\triangle r(t)$ might have some strong smoothness properties that would provide great robustness to the process. But in the absence of information about the smoothness of the relative range profile, a reasonable algorithm is to pick the range for the next sample that is closest to the estimate of the range of the current sample. This is equivalent to assuming that the phase change between samples is small. This leads to a simple algorithm that can be captured 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.

3.3.2 Unwrapping Errors

There are two sources of error in the resulting range profile: 1) phase noise, and 2) phase unwrapping errors.

Phase noise error is easy to understand; it is η_i for each sample. Phase unwrapping errors are caused by a sequence of phase noise, for example, a large positive error followed by a large negative error, that cause a jump of ± 1 in the sequence of k_i's. This is illustrated in Figure 3.

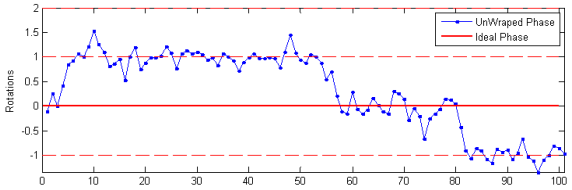

Figure 3. Phase noise errors and three unwrapping errors.

Notice that phase noise errors have no cumulative effect in the unwrapped phase. A positive noise spike on a given sample will cause the phase change for that sample to be large and positive and the phase change on the next sample to large and negative such that the results cancel out. The phase noise portion of the error in the unwrapped results depends only on the first and the last sample.

However, phase unwrapping errors persist forever (or until a new phase unwrapping error cancels it out), as is seen in Figure 3.

Instead of errors accumulating like when summing a long sequence of noise measurements, errors are corrected on the next sample, except that every once in while when a phase unwrapping error occurs, at which point the relative range profile jumps by $\pm\lambda/2$. In intuitive terms, it is as if the "erroneousness" of the measurement process is suppressed for long periods of time and then realized in quantized bursts which are always $\pm\lambda/2$ in magnitude. This is actually quite a useful characteristic of the relative range profiles produced by phase unwrapping, but it is important to avoid methods that depend on the error having a Gaussian distribution.

3.3.3 Measuring the Accuracy of the Relative Range Profiles

We wished to test the accuracy of the range profiles produced by phase unwrapping on this hardware. In order to do this we setup a small pendulum shown in Figure 4. Because the pendulum produces well understood motion profiles, with only a few unknown parameters (that define the starting position and the drag coefficient), it allows us to measure both the phase noise induced errors and the rate at which phase unwrapping error occurred. However, as pointed out in the last section, only the rate of phase unwrapping errors affects displacement detection performance.

We found that the phase noise portion of the relative range profile error was about 5 mm and that phase unwrapping errors occurred about once every 10 to 15 seconds, on average, depending on environmental noise levels. If the likelihood of positive errors is the same as the likelihood of negative

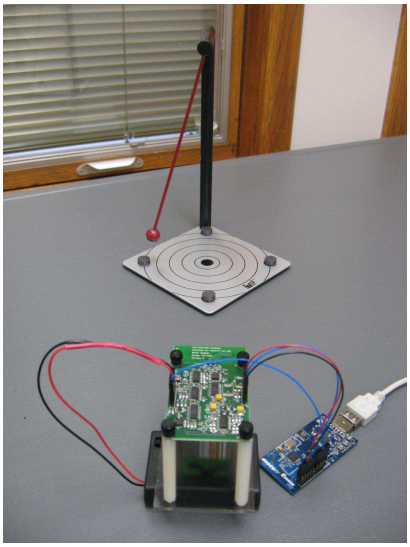

Figure 4. The setup used to measure the accuracy of the relative range profiles produced by phase unwrapping the output of the BumbleBee.

Table 2. Extrapolated growth of errors.

$\triangle t$	Error
10 s	2.6 cm
1 min	6.3 cm
10 min	20 cm
1 hr	49 cm
4 hr	98 cm
1 day	2.4 m

errors then the accumulation of errors would grow proportional to $\sqrt{\triangle t}$ for large $\triangle t$. This asymptotic result is shown in Table 2.

3.3.4 Caveats

These results show that our relative range profiles are remarkably accurate for mote-scale sensors. In Section 4.1 we carefully measure the accumulation of phase unwrapping errors in operational environments. The results of **Section 4.1** are not strictly comparable with **Table 2**, because in Section 4.1 we are concerned with the one in 10 thousand or one in a million worst accumulations of error, rather than the average accumulation of error. In addition, we discovered that the complexity of moving bushes caused more errors than the well behaved motion of the pendulum.

Finally, in some settings, we observed narrow-band interference, from unknown sources, that caused substantial phase drift, even in the absence of phase unwrapping errors. This problem only manifests itself in very low return environments (i.e., those devoid of any moving reflectors); the presences of a real return swamps the effects of very weak interferers. As a result, this problem can be detected and avoided in software; if the return is much smaller than the return from a human, phase-unwrapping to decide if you have an intruder may perform poorly.

4 Experimental Results

In the first part of this section, we experimentally measure the quality of phase-based net displacement measurements in realistic outdoor environments. In order to validate the FAR for the associated Displacement Detector, it is necessary to estimate the probability distribution of the net displacement both for environments with and without an intruder. In particular, we are interested in frequency or severity of statistically unlikely events; the analogy is estimating the severity of the 50 year flood or the 500 year flood so that we know how high to build the levees. Specifically we need to estimate the expected worst case accumulation of phase errors over any 3 second window in a week or a year.

The second part of this section presents implementation details for implementing the detector on an actual mote.

4.1 False Alarm Rate Validation

To estimate the FAR we performed long running data collection experiments in a variety of environmental settings with different vegetation, wind and background clutter conditions. This includes data collected in the presence of bushes (small and large), trees, on paved surfaces, in public parking garages, in office environments with other wireless interferers, in adverse weather such as during snowfall and light rain, to but name a few. We also collected long running traces (sometimes for days or weeks), using a video camera to verify that there were no unexpected intruders during the data collection.

In order to better understand the PD we also collected 100's of data sets with a human walking at carefully measured distances.

In this subsection, we focus our results to the analysis of data collected in four different environments: 1) an empty, mowed open field with no moving background clutter, 2) in a public parking garage with stationary cars parked nearby, 3) near a group of small bushes under lightly windy conditions, and 4) near a group of tall trees under highly windy conditions.

4.1.1 Comparing Amplitude with Phase

Figure 5 shows the radar output for the four environments described above over 30 minute periods. It is common to talk about *noise* as the random (or perhaps nearly random) measurement error, *clutter* as the return the environment, and *signal* as the return from targets. At times when talking about signal processing, the term signal is also used to mean the output of the radar, which would be the sum of all three of these "signal-components", but in this section we will refer to signal only as the return from targets. The key point is that none of these graphs contain any signal, they are pure clutter.

The other key point is that even under light wind conditions or occasional wind gusts, moving background objects such as bushes or trees can result in clutter levels that are comparable to, or actually exceed, the expected signal levels.

Figure 6 shows the variation of the amplitude and the unwrapped phase as a function of target range. The graph on the top shows how the amplitude of the radar return changes as a target moves from 9 m to 2 m. As expected, the amplitude of the return is a function of the target size (or more

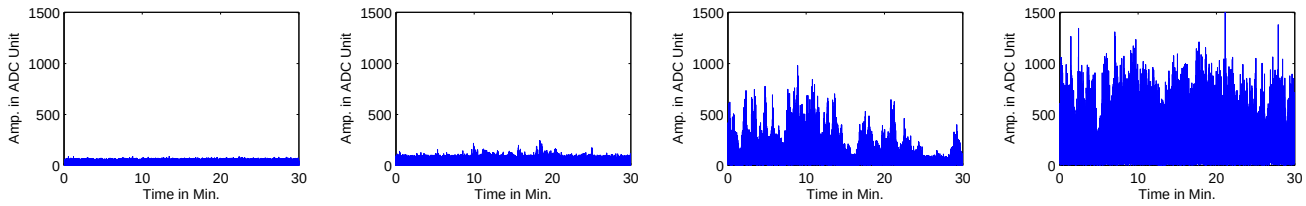

(a) Empty mowed field (b) Public parking garage (c) Small bushes in light wind (d) Large trees in heavy wind

Figure 5. The amplitude of the radar output over over 30 minute periods.

precisely, its Radar Cross Section (RCS)) and the distance from the sensor. In Figure 6, the red line shows the change in the unwrapped phase measurement for a 1 m change in the target's location. The fact that it is nearly constant at 1 m indicates good performance. The blue line shows the relative range profile as a function of actual range.

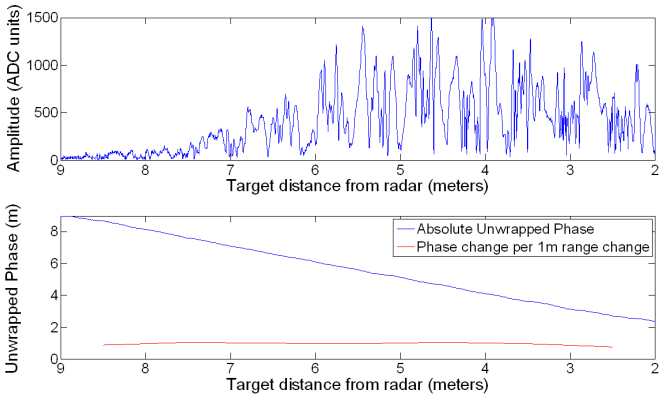

Figure 6. Amplitude and unwrapped phase measurements for the same target walking towards the radar. The flatness of the red line indicates that the measured displacement is independent of target range.

There are two interesting features of the unwrapped phase. First, the accuracy of the unwrapped phase seems to only be weakly affected by the signal level. Secondly, for very low SNRs, e.g., in the vicinity of 9 m range, the data still produces good relative range profiles. In fact, even in the absence of TNIs, when the amplitude based detection performs the best, the phase-based indications of motion work at lower SNRs than the amplitude based indications of motion.

4.1.2 Modeling the Unwrapped Phase Distribution

In order to estimate the height of the 500 year flood it is not usually practical to collect 500 years of data. Rather, the normal approach to this problem is to collect a large amount of data, model the probability distribution, and extrapolate into the tail of the probability distribution.

Estimating and extrapolating a probability distribution carries the risk that you use the wrong model for the probability distribution, but there is really no alternative to estimate very rare events. In order to estimate FARs that are less than about once a week, we are forced to use this method. We expect that the distribution of the sum of all the phase unwrapping errors over an arbitrary short window (like a 3 second window) might be Gaussian. Even over a 3 second window, there are about about one thousand independent noise events that contribute directly to potential unwrapping errors, so it seems that the Central Limit Theorem might apply. However, there might be an extensive network of correlations between phase measurements from a bush swaying in the wind, and short of experimental validation it is hard to be sure that the distribution should be Gaussian.

Figure 7 shows the cumulative distribution function, on a semilog scale, of the net phase change (translated into real distance units), over a 3 second window for the clutter (i.e., no target) cases shown in Figure 5. A quadratic fit (on a semilog scale) fits the data particularly well, so we feel somewhat confident in extrapolating with a Gaussian distribution. We also performed a smaller number of much longer data collections to further validate this model.

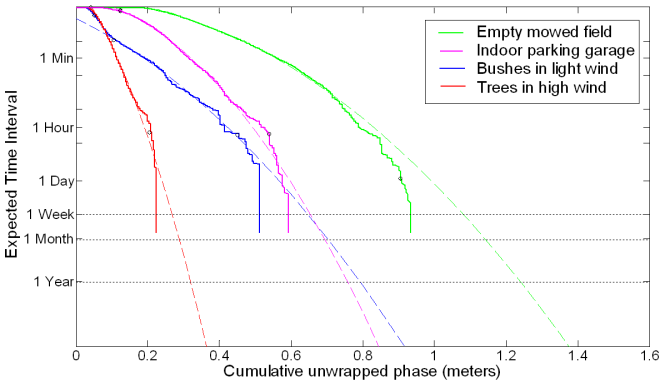

Figure 7. The cumulative probability distribution for unwrapped phase over 3 second windows for clutter (i.e., non-target) data shown in in Figure 5.

4.1.3 False Alarm Rate Results

This extrapolation indicates that the 1 week worst case event is a combination of noise that results in enough phase unwrapping errors to cause 1.08 m of error accumulation over a 3 second window. This is a shockingly bad event when you consider that the results in Section 3.3.3 suggest on average a 3 second window probably doesn't even have a single unwrapping error, but every once in a while really bad things happen. However, this is still a very good result because it suggests that the detection threshold can be as low

as 1.08 m, which will yield high PD, is sufficient to achieve a very good FAR.

Interestingly, the best case for the unwrapped phase model is actually the one for which the clutter amplitude is the highest; this corresponds to the highest Clutter to Noise Ratio (CNR). In the most "benign" environment, i.e., where there is no moving object, the CNR is very low and hence there are a lot of unwrapping errors.

This worst case could be identified and alleviated by using a hybrid algorithm, where phase unwrapping is performed only when the amplitude of the radar output exceeds a minimum level. This hybrid algorithm yields much better performance, but in order to make the analysis of the performance more reliable we will only analyze the pure phase-based algorithm.

4.2 Overall Performance Evaluation

4.2.1 Probability of Detection

In order to measure the PD we had a human walk through the scene several hundred times, so that the Closest Point of Approach (CPA) was 7 m from the radar. Figure 8 shows the cumulative distribution of the cumulative phase change over the same 3 second integration windows. The blue curve shows the distribution when walking directly (radially) towards the sensor sensor stopping at a range of 7 m and the red curve shows the distribution when the target walks past the sensor (orthogonally) in a straight line such that the CPA is at 7 m. Notice that the PD is essentially 1 up to a FAR of 1 every 84 days, and that it is 0.97 for 1 false alarm per year.

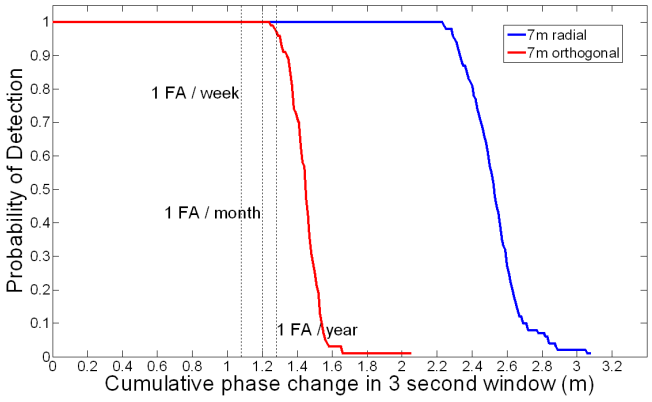

Figure 8. The cumulative probability distribution for unwrapped phase over a 3 second window for a human with a CPA of 7m.

In many practical deployments where more than one sensor is expected to detect a target , much lower node level PDs would yield very high network-level PD. In this case it would be reasonable to raise the threshold to obtain even lower FARs.

4.2.2 Amplitude Based Detection

We used the classical amplitude based M-out-of-N detection algorithm which has been extensively used in theory [22] and in practice [23,24] for intrusion detection using radar as well as other sensors such as PIR and magnetometer. We used data from several hundred runs to calculate the amplitude returns from the radar for a human target at various distances. Based on this distribution, we carefully selected the M-out-of-N parameters to optimize detection performance with respect to both probability of detection and false alarms.

From Figure 5, we observe that the noise amplitude and hence the SNR for a real target can vary significantly depending on the environment. We now compare the detection performance of Phase Unwrapping based Displacement Detection with the classical M-out-of-N amplitude based detector. Figure 9 shows the probability of detection and the corresponding rate of false alarms under various settings.

4.2.3 Overall Comparison

It should be noted that the amplitude based detector has extremely small FAR in the parking garage and in the empty field where the background clutter is stationary. However, as seen from Figure 9, for a detection range of 7 m and a detection probability of 0.5, the amplitude based detector produces about 1 false alarm per minute under light wind conditions, which degrades to 1 false alarm every 9 seconds when placed near trees blowing in strong winds.

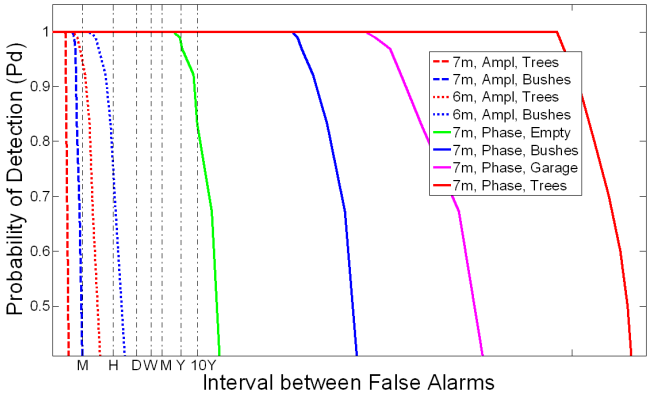

Figure 9. PD vs FAR for Displacement and Amplitude Detectors for different detection range and environment settings. The X-axis is shown in terms of false alarm interval for easier understanding. The amplitude detector is omitted for the garage and the empty field, because in these environments it has negligible FAR.

By contrast, the Displacement Detection algorithm, for a 7 m walk by, produces 1 false alarm every 84 days with a detection probability of essentially 1 in the worst case. And it achieves 1 false alarm per year with a detection probability of 97%. Moreover, this performance is achieved without compensating for the phase drift in low CNR environments.

4.3 Mote Implementation

The BumbleBee radar produces two outputs that form the real and imaginary parts of the complex radar returns. In order to calculate the unwrapped phase, we first need to calculate the "wrapped" phase for each complex sample. If we denote the complex data samples as C_k, then $\phi_k = \tan^{-1}(\Im(C_k)/\Re(C_k))$. The arc-tangent function is not

easy to implement on a mote. In this section, we describe two implementation options for resource constrained motes.

4.3.1 Phase Lookup Table

Mote platforms are typically constrained in the amount of RAM available and the amount of processing that can be performed per sample. However, motes typically have ample program memory where lookup tables can be stored by declaring them at compile time. We used this approach to implement a lookup table for $\tan^{-1}$.

In theory, $\Im(C_k)/\Re(C_k)$ can span the entire range of real numbers, however using basic trigonometric principles, we can restrict our lookup table to the range $(0,1)$. We create a table of size N by computing $tan^{-1}(x)$ offline for $x = \{0/N, 1/N, 2/N, \ldots, N/N\}$. Then during phase unwrapping, we can calculate the wrapped phase by looking up the ratio of some combination of $\Im(C_k)$ and $\Re(C_k)$ that lies in the range $[0,1]$. For example if $\Im(C_k) > 0$, $\Re(C_k) > 0$, and $\Im(C_k) > \Re(C_k)$, we use the property that

$$tan^{-1}(x) = \pi/2 - tan^{-1}(1/x)$$

to lookup $\Re(C_k)/\Im(C_k))$ instead.

4.3.2 Calculating the Number of Rotations

In the phase unwrapping algorithm described above, we calculate the unwrapped phase for each sample. However, for Displacement Detection, we only need to know whether the phase difference in a window of time exceeds some threshold. Thus, although the physical understanding of the algorithm is that it computes the phase, unwraps the phase, and then compares the change in phase to a threshold, if all we care about is the final answer, i.e., whether the change in phase exceeds a threshold, then the computation of the phase is an intermediate result that may or may not be necessary. In this case, it is possible to compute the final answer without ever computing the phase. The approach is to estimate the cumulative number of phase wrappings or rotations (which can be easily translated into phase by multiplying by 2π). For T$\gg$1 rotation, this can be approximated using integer operations as follows:

$$\begin{aligned}
Rot(k) &= Rot(k-1)+1 \quad if \quad \Im(C_k)*\Re(C_{k-1}) > \Im(C_{k-1})*\Re(C_k) \\
&\qquad \&\&\ \Im(C_k) > 0 \ \ \&\&\ \Im(C_{k-1}) < 0 \\
Rot(k) &= Rot(k-1)-1 \ if \quad \Im(C_k)*\Re(C_{k-1}) < \Im(C_{k-1})*\Re(C_k) \\
&\qquad \&\&\ \Im(C_k) < 0 \ \ \&\&\ \Im(C_{k-1}) > 0 \\
Rot(k) &= Rot(k-1) \quad otherwise
\end{aligned}$$

Calculating the cumulative number of rotations in this manner requires only two integer multiplications and a few comparison operations.

4.3.3 Converting Phase/Rotations to Real Distance

The unwrapped phase or number of rotations is directly correlated to the actual physical distance traversed by the moving target in the radial direction. This conversion is useful while setting thresholds which are based on target motion. Since the BumbleBee radar operates at a 5.8GHz center frequency, the wavelength (λ) of the radar is 5.17 cm, which corresponds to two phase rotations or 4π radians. Thus, we can calculate the radial distance from the unwrapped phase or cumulative rotations as

$$\triangle r_k = \phi_k \cdot \frac{\lambda}{4\pi} = Rot(k) \cdot \frac{\lambda}{2}.$$

Thus, one phase rotation corresponds to a physical radial distance of 2.58 cm. From the analysis and experimental results in Section 4, we find that in practice, the threshold for Displacement Detection is greater than 1 m. This threshold is thus $\gg 1$ rotation, which makes the accuracy of the rotation approximation algorithm acceptable.

4.3.4 Calculating Cumulative Phase Change

We declare a Displacement Detection if the cumulative phase change over an integration time window exceeds some threshold. To calculate this change, we need to compute the running minimum and maximum of a circular buffer of the size of the integration window.

We use a classical heap-based algorithm to compute the minimum and maximum. We use the trick of having two index pointers, one from the heap order to the circular buffer order and a reverse index pointing from the circular buffer order to the heap order. Using this bi-indexed structure, we only need $O(log(N))$ time to update the data structure upon replacing the oldest entry with the new one. Specifically, we require (log_2(N) - 1) swaps in each update. For a window size of 3 seconds at 341Hz, this translates to roughly 10 swaps per sample on average, which can be easily accomplished on a mote.

5 Conclusions and Future Work

In this paper, we showed that phase information from medium bandwidth mote-scale radars can produce relative range profiles with relative accuracy that are as good as the range profiles produced from much higher power UWB radars. We further showed that a very simple algorithm can utilize this information to perform Displacement Detection, which effectively eliminates false alarms caused by vegetation blowing in the wind.

We collected long data samples in a variety of environments and constructed probability models from this data. Extrapolation showed that our simplest Displacement Detection algorithm could reduce the component of the FAR due to moving vegetation to about once every three months in the most challenging environments. In fact, the extrapolation suggests that the component of the FAR due to moving vegetation would be below once per year in most environments; but more data is probably required to fully support this claim.

The phase unwrapping and Displacement Detection algorithms described in this paper can be very low power. We are currently using a 40 mW COTS radar, but future generations of the hardware could lower the radar power by an order of magnitude. More critically, we are able to implement the signal processing software in less than 5 mW on a COTS mote. Finally we have implemented a version of this Displacement Detection algorithm on a low-power FPGA in less than 100 μW.

The relative range profiles can also be used for many applications other than Displacement Detection. Profiles from

multiple radars can be fused to achieve high precision in-network target tracking. The relative range profiles can also be used to estimate velocity. Tracking the phase of the return is also useful in recognizing subtle patterns of target motion and has applications in monitoring activity levels, analyzing the motion of rotating machinery, etc.

In addition, the complex output of the radar contains other fine scale motion information not fully represented in the relative range profile that offers the potential to solve other TNIs problems. Examples include distinguishing between reptiles, bipeds, and quadrupeds. Such approaches are closely related to MicroDoppler analysis, [25].

We found that instrumenting ground truth for data collection was an invaluable tool while analyzing data collected across different times and places. We augmented our data collection software with camera for capturing details of background and target motion and weather sensors. Both the data and the software library and hardware design tools used in the data collection are available for sharing with the research community.

6 References

[1] P. Dutta, A. Arora, and S. Bibyk. Towards radar-enabled sensor networks. In *Proceedings of the 5th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 467–474, 2006.

[2] AFRL. 2009 commander's challenge. `http://www.wpafb.af.mil/news/story.asp?id=123196732` .

[3] T. He et al. Vigilnet: An integrated sensor network system for engery-efficent surveillance. *ACM Transaction of Sensor Networks (TOSN)*, 2(1):1–38, Febuary 2006.

[4] K. Lin et al. Heliomote: enabling long-lived sensor networks through solar energy harvesting. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 309–309. ACM, 2005.

[5] P. Dutta et al. 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.

[6] J. Yick, B. Mukherjee, and D. Ghosal. Wireless sensor network survey. *Computer Networks*, 52(12):2292–2330, 2008.

[7] B. Flanagan and K. W. Parker. Robust distributed detection using low power acoustic sensors. *Procedings of SPIE*, 2005.

[8] A. Quanch and K. Lo. Automatic target detection using a ground based passive acoustic sensor. *Conference on Information, Decision and Control*, pages 187–192, 1999.

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

[10] J. Cihlar, T. J. Pultz, and A. L. Gray. Change detection with synthetic aperture radar. *International Journal of Remote Sensing*, 13(3):401–414, February 1992.

[11] R. G. White. Change detection in sar imagery. *International Journal of Remote Sensing*, 12(1):229–260, February 1991.

[12] L. M. Wells and A. W. Doerry. Synthetic aperture radar: Not just a sensor of last resort. *Proceedings of SPIE*.

[13] Time Domain. Pulson220 datasheet. `http://www.timedomain.com/products/P220aRD.pdf` .

[14] Fcc 02-48. `http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-02-48A1.pdf` .

[15] G. S. Monk. An adaptation of the interferometer for hyperfine structure. *Review of Scientific Instruments*, 4(10):527–528, October 1933.

[16] C. J. Hubbard and I. F. Zartman. A fixed path acoustic interferometer for the study of matter. *Review of Scientif Interments*.

[17] L. R. Hafstad and M. A. Tuve. An echo interference method for the study of radio wave paths. *Proceedings of the IRE*, 17(10):1786–1792, October 1929.

[18] Radar interferometry. *IEEE RADAR '08*, May 2008.

[19] G. Schultze B. Koch H.D. Vom Stein, P. Rateau. New laser interferometry methods of measuring the velocity of high-speed model missiles. *Radio and Electronic Engineer*, 40(1):45–48, July 1970.

[20] M. Maróti et al. Radio interferometric geolocation. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 1–12. ACM, 2005.

[21] I. Amundson, J. Sallais, X. Koutsoukos, and A. Ledeczi. Radio interferometric angle of arrival estimation. In *7th European Conference on Wireless Sensor Networks*, Coimbra, Portugal, 02/2010 2010. Springer, Springer.

[22] J Toomay and P Hannen. *Radar Principles for the Non-Specialist*. SciTech Publishing, 3rd ed, 2004.

[23] L.Gu et al. Lightweight detection and classification for wireless sensor networks in realistic environments. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 205–217, New York, NY, USA, 2005. ACM.

[24] A. Arora et al. A line in the sand: A wireless sensor network for target detection, classification, and tracking. *Computer Networks, Special Issue on Military Communications Systems and Technologies*, 46(5):605–634, July 2004.

[25] M. G. Anderson. Design of multiple frequency continuous wave radar hardware and micro-doppler based detection and classification algorithms. 2008.

[26] K. J. Hintz. Snr improvements in niitek ground-penetrating radar. *Detection and Remediation Technologies for Mines and Minelike Targets IX*, 5415(1):399–408, 2004.

APPENDIX

A Using Radar Phase Information

A.1 Pulsed Doppler Radars

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

A.1.1 Continuous versus Pulsed Wave

The difference between traditional Doppler Radars and a 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 component of the return which is associated with moving targets. This type of PDR requires that the returns be coherent between pulses. That is, an exactly identical environment needs to produce the same result from one pulse to the next. Coherence also implies that the SNR 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 is not coherent (which is typically the case). Nevertheless, real electronic systems cannot maintain perfect coherence indefinitely and the limit to which coherence can be maintained establishes a lower bound on the Doppler frequency which can be estimated. We describe next the COTS hardware we used that that achieves coherence over 2 or 3 seconds.

A.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 10(a). The correlation of this pulse with itself is shown in Figure 10(b). The key feature of the system response of this radar 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.

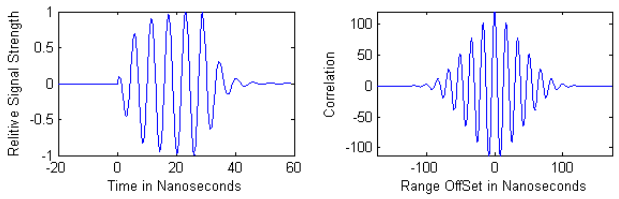

(a) (b)

Figure 10. A prototypical pulse for a PDR (a) and its autocorrelation (b).

(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 [26], 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.)

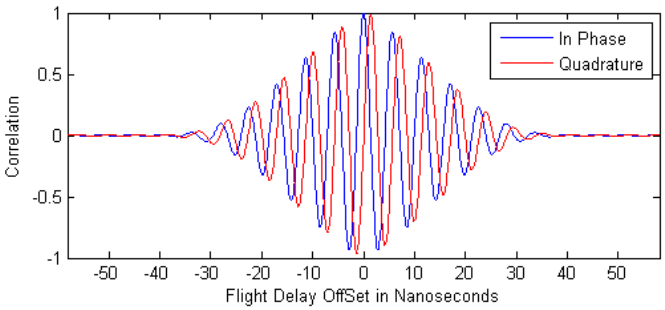

Figure 11. Output response for a "complex output" PDR.

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 Figure 10 are shown in Figure 11. 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; one may also represent the sample in terms of its *amplitude* (aka absolute value/modulus) and *phase* (aka argument/angle).

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 12.

B Phase Information

In this section, we first motivate why amplitude information from the complex output PDR described above is insuf-

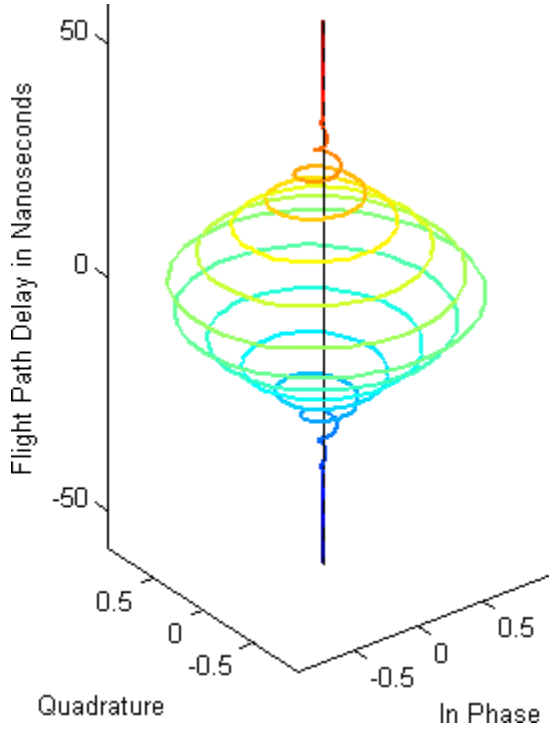

Figure 12. The display of Figure 11 as a complex real valued function.

ficient 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. Finally, we 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.

B.1 Amplitude vs. Phase of Complex Output

Notice from Figure 12 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.

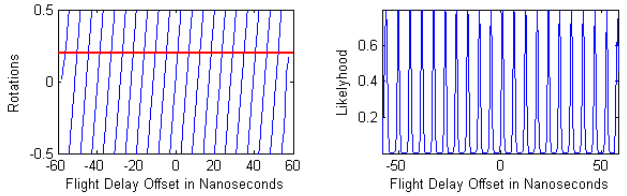

(a) (b)

Figure 13. Phase (expressed in units of rotations) as a function of the range to the target(a) and the likelihood function for the data in (a) for an SNR of 3dB(b).

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 13(a) 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 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 13(b) 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 13(b) 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 is a function of SNR. 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.

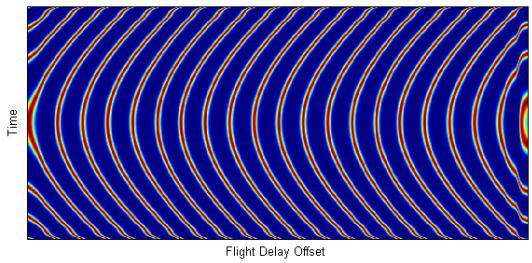

Figure 14. The evolutions of the likelihood function within a range bin as a target passes through the field of view. The red lines corresponds to maximum likelihood and the blue region to near-zero likelihoods.

Figure 14 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 target's relative motion to an accuracy that is a fraction of the wavelength, even when the range bin is 10 to 100 times larger.

In sum, we claim that with phase measurements, it is possible to obtain *locally precise but globally ambiguous* range information.

B.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 easily seen by the human eye in Figure 14. 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 realizing phase unwrapping, we have to deal with discontinuities that occur in the measured phase, see Figure 13. These discontinuities arise from a discontinuity in the phase function. To better appreciate this problem, it is helpful to realize that our interest in the phase is essentially to measure the amount of rotation that has occurred from an arbitrary starting point (see Figure 12). The fundamental problem is that one could traverse between any two points in the complex plane in a clockwise direction or 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 15. The cut is along the negative real axis; this choice of cut is arbitrary, but a cut is always required.

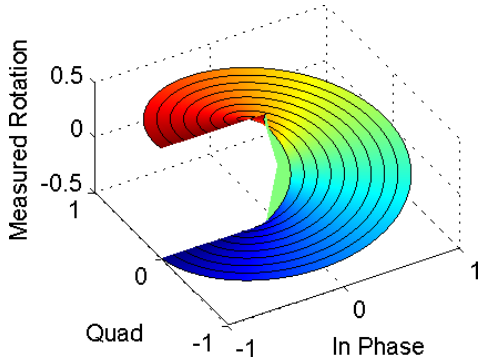

Figure 15. The standard cut in the phase function.

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 well-known process of constructing a trajectory on the Riemann surface shown in 16.

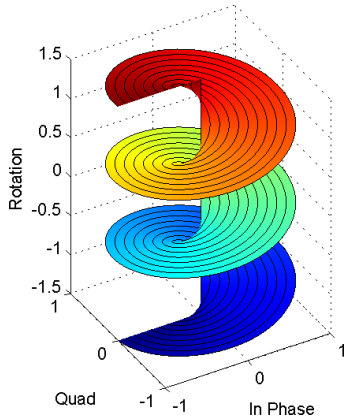

Figure 16. The Riemann surface that is the extension of the phase function shown in Figure 15.

B.3 Tracking using Phase Unwrapping

The basic phase unwrapping algorithm implicitly exploits a random walk style motion model, which assumes that the next location is close to the old location. For targets with significant momentum, it would be better 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 that is closest to the predicted value and corresponds to the measured wrapped phase.

For instance, when applying phase tracking to track a 1-dimensional pendulum in this manner, we use a quadratic motion model. We 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 17.

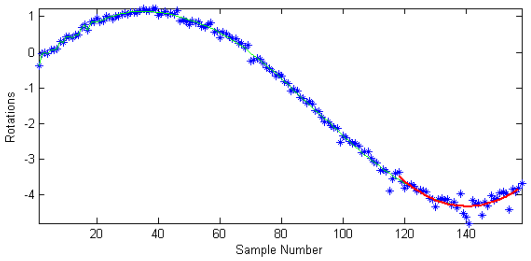

Figure 17. An example execution of partial phase tracking. The blue points are the unwrapped phase, the green line is the estimated track, and the red line is the quadratic motion model that fits the last 40 points.

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.