Automatic Animal Tracking Using Matched Filters and Time Difference of Arrival

Robert MacCurdy
Cornell University/Mechanical and Aerospace Engineering, Ithaca, USA
Email: rbm7@cornell.edu

Rich Gabrielson, Eric Spaulding, Alejandro Purgue, Kathryn Cortopassi,
Cornell Laboratory of Ornithology/ Bioacoustics Research Program, Ithaca, USA

Kurt Fristrup
National Parks Service, Fort Collins, USA

***Abstract*—A method for tracking animals using a terrestrial system similar to GPS is presented. This system enables simultaneous tracking of thousands of animals with transmitters that are lighter, longer lasting, more accurate and cheaper than other automatic positioning tags. The technical details of this system are discussed and the results of a prototype are shown.**

***IndexTerms*—radio tracking, TDOA, wildlife tracking, CDMA, localization, pseudolite, GPS, matched filter**

I. INTRODUCTION

Radio-tracking has been widely used to monitor wildlife movements since the 1960s [1,2] with hundreds of scientific papers published using some variant of this method. For the majority of these studies, an operator monitors received signal strength while changing the orientation of a directional receiving antenna. The direction yielding the maximal signal strength is recorded as a pointing vector to the tagged animal. This simple method is adequate to guide a researcher to the location of a focal individual, and triangulation using two or more receiving stations can be used to track a few individuals simultaneously. The accuracy of this method is detailed in [3]. However, this method yields relatively few position fixes per hour, and fully absorbs an operator's attention and effort. Automatic tracking systems have been developed using fixed receiving towers [4], [5], [6]. Most efforts involve directional antennas and rely on the beam pattern of the antennas to infer a direction of arrival. These approaches generally suffer an error in the 1-10 degree range, depending on the implementation [6] and the cross-bearing positional error for each receiving station increases linearly with range.

In addition to terrestrial radio-tracking methods, two satellite based options exist: GPS and Argos. These two systems provide location information using different techniques. GPS employs a network of orbiting satellites that broadcast signals to a terrestrial receiver which uses a Time Difference of Arrival (TDOA) algorithm to estimate its position. The Argos system exploits the frequency shift in the received transmission (Doppler shift) caused by a satellite's motion relative to a transmitter on the earth. The principle drawback of GPS is that it does not directly provide a means of reporting position information back to the researcher. The position information is either stored and retrieved later, or downloaded via an auxiliary radio frequency link. Weight is also a limiting factor for satellite-based systems. The typical maximum allowable tag to body weight ratio is 5%, with lower being preferable. The smallest GPS tags at present are in the 22 to 150 gram range, which limits their application to larger animals (>440 g). The Argos system can achieve reasonably good accuracy; the best service available advertises 150 m accuracy. However, this level of accuracy is not often available, and the other three accuracy classes range from 150 to 1000 m. Argos tags are generally smaller than GPS tags, but still only allow tracking of animals heavier than 200 g. This weight constraint excludes 40% of all bird species [7]. The cost of GPS and Argos tags is also prohibitive for large-scale studies. The complexity and low volume of these tags lead to typical single unit prices in the $1500 to $4000 range, with little cost reduction at larger volumes.

The weight, cost, and performance of existing radio-tracking techniques necessitate a new approach. We have designed, built and installed a prototype system based on Code Division Multiple Access (CDMA) and Time Difference of Arrival (TDOA) that is capable of automatically locating thousands of tags in real time. This system borrows concepts from the GPS system in general, and from pseudolites in particular. Pseudolites (or pseudo-satellites) refer to terrestrial devices which transmit GPS signals. An excellent summary can be found in [8]. Pseudolites have existed for many years in various forms; in fact some of the earliest pseudolite tests actually predate the launch of GPS satellites [9].

II. SYSTEM DESIGN

A. System Overview

Previous attempts at terrestrial automatic wildlife radio tracking have relied on conventional, fixed frequency, narrow band transmissions. These transmissions carry relatively little information for the power consumed during transmission. In contrast, our approach relies on a new, broadband transmitter that emits a transmission

Manuscript received December 14, 2008; revised May 12, 2009; accepted June 15, 2009.

modulated with a unique pseudo-noise (PN) sequence. This approach is similar in many ways to a GPS system operating solely in acquisition mode, however in our system, the mobile device to be tracked is the transmitter, rather than the receiver. The transmitted sequence occupies a typical bandwidth of 2MHz and lasts about 3 msec. Each transmitter operates on the same carrier frequency, which is currently 140MHz, but each transmitter modulates that common carrier with a different PN code. We have chosen Gold Codes for the system, due to their attractive cross-correlation properties. This yields a code division, multiple access system: many transmitters can operate simultaneously without significant interference. The receiver, unlike a conventional animal tracking receiver, "listens" continuously to the 2MHz wide bandwidth that the transmitted signal occupies and runs a matched filter detector. When a signal is detected, the receiver precisely timestamps the event and sends a data packet to a server with the event ID, tag ID, time and other information. The server accumulates these individual data packets from different receivers, which are located in different positions, and groups receive events. When multiple receive events correspond to the same transmission event (the tag was "heard" at multiple receivers), the server computes a position estimate. Three or more receivers must receive the same transmission in order for the server to compute an accurate position fix. If only two receivers hear the transmission, the transmitter position can be limited to a particular hyperbola, and if only one receiver hears the transmission, a simple presence/absence data point is available.

The system uses precise timing of receive events to calculate transmitter position. One major benefit of this approach is that the accuracy of the system does not degrade as a direct function of distance, unlike the Direction Of Arrival (DOA) based systems. Significant power savings, relative to conventional techniques, are also realized by this approach. Very short and infrequent transmissions are possible because the receivers are listening constantly, and can listen for transmissions from all tags simultaneously. Conventional radio tracking tags transmit for 30 msec and repeat every 2 to 3 seconds. This transmission length and interval are the minimum allowable times for a human operator to accurately discern a reception event and determine tag direction. However, in our system the tag can be configured to transmit for only 3 msec and repeat only as often as independent position fixes are required. For example, it is entirely possible to transmit once every 5 min, which reduces tag power consumption by a factor of 1500, relative to traditional tags. This dramatic reduction in consumption enables multi year studies with transmitters similar in mass to conventional tags. If maximum tag lifetime is not a priority, lower power consumption can also enable lighter tags. Our tags are easily programmed, in contrast to conventional radio tracking tags, which must be custom-built for a particular frequency and cannot be readily modified.

B. *The Transmitter*

The transmitter (tag), shown in Figure 1, is based on an inexpensive, very low power microcontroller, along with a separate PLL, mixer and amplifier. Our design integrates off-the-shelf ICs in order to avoid the high cost and long development time of a custom ASIC. This choice results in an implementation that is larger than it could otherwise be, but this tradeoff allows rapid development. The tag's 140MHz center frequency implies a ¼-wave antenna length of approximately ½ meter. This is too long for most small birds to manage, so the actual antenna used is often between 15 and 25 cm. Despite the efficiency penalty that these electrically short whip antennas impose; they are very common in animal tracking applications. The tag uses a binary phase shift keyed (BPSK) modulation scheme to directly modulate and spread the carrier power, as shown in Figure 2. The modulation rate is 1MHz, resulting in a 2MHz wide main band. The tag is programmable for center frequency, transmission interval, pseudo-noise code, chip-rate, RF output power, and operating schedule. This programmability allows tailoring the tag parameters to the application, which maximizes lifetime for a particular tag mass. The current tag mass is 1.4g without the battery and encapsulation. Future versions will reduce mass by 50%.

Transmitting an RF signal is the most power consuming operation for any small transmitter. In conventional tracking systems, many of these transmissions are not used: receivers can only tune one tag at a time, human operators require multiple transmissions per bearing, and multiple bearings are required to determine a location. In contrast, the automatic detection algorithm running in our receivers is capable of determining tag location from a single transmission. This capability enables a very low duty cycle, dramatically increasing tag lifetime.

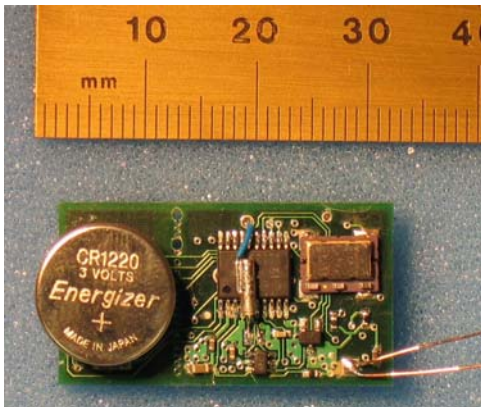

Fig. 1 Transmitter board, with typical battery and programming tab still attached

Each tag sends two PN sequences, one after the other, per transmission interval. The first sequence is common to all the tags and serves as an acquisition and synchronization signal. The second sequence is unique to each tag and encodes ID. Each sequence is chosen from the family of Gold [10] codes available from an 11 bit generator. Gold codes are attractive due to their predictable cross-correlation performance and relatively easy generation. The choice to send two PN sequences was motivated by the limited processing power of the receiver stations. In principle, each tag could send only its unique ID sequence. This approach would require the receivers to search in code ID space as well as in code time-alignment. The DSP processing power available is insufficient for this task. Therefore, we employ a two stage detection process: search continuously and in real time for the acquisition and synchronization signal, and when it is detected, search in code ID space to identify the tag. This latter step need not take place in real time. This approach, and our current implementation, causes the system to block subsequent tag detections during the 3 milliseconds that the signal from the first detected tag occupies. This reduces the multiple access potential that true CDMA techniques offer. We reduce the likelihood that tags will repeatedly interfere with each other by introducing a random offset into their transmit intervals. If two tags interfere during one interval, they are very unlikely to interfere in the near future. Very short tag transmit durations also mitigate interference. A system employing all 2049 possible IDs, with each tag transmitting every minute, would only be using 10% of the available time. Systems with fewer tags or longer transmit intervals would experience reduced interference.

The family of Gold codes we use supports $2^{11}+1$ individual IDs operating simultaneously in a particular study area.

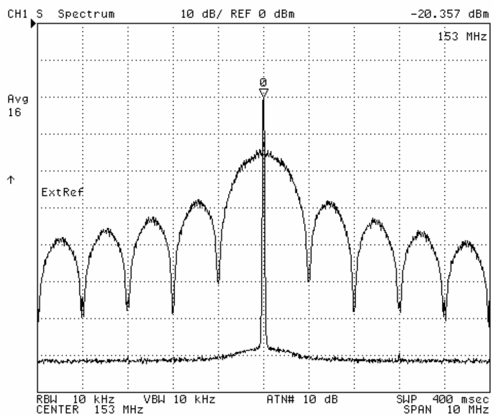

Fig. 2 Transmitter output showing unmodulated and modulated carrier

The system uses 2049 unique IDs, but identical tag IDs could potentially be assigned to more than one individual if the expected individual movement patterns and continuity of tracking are likely to provide decisive identification. For example, the same ID could be given to a tag on a rabbit and a tag on a wolverine, and there would be very little chance of confusion. A similar argument could be applied to animals whose usage of an area is clearly demarcated by season.

The population of tags tracked by this system can be orders of magnitude larger than systems that utilize carrier frequency to identify individuals. Automatic tracking with narrow band systems does not scale well for two reasons: the available bandwidth will be rapidly consumed by tags whose crystal tolerances require at least several kilohertz of spacing, and the scan speed must be increased as the scanned bandwidth increases in order to avoid missing a transmission event. The latter requirement places a hard lower limit on the duration of the transmitted signal from a traditional tag, though future scanning receivers could address this problem by abandoning the traditional super-het topology in favor of a wideband, software defined architecture. These considerations indicate that both RF spectrum allocation and engineering effort will be more efficiently translated into tracking capacity using the CDMA approach.

C. *The Receiver*

Figure 3 shows a single receiver system, including the equipment box, tripod, batteries and antenna mast, which extends 5 meters above what is shown in the image. The receiver uses a 2-7/8λ phased element monopole antenna that yields approximately 6dB of gain. The system requires approximately 10 Watts and at 20 kg (not including lead acid batteries) can be easily moved.

The receivers use a method similar to GPS in order to automatically locate animals wearing our tags. Each receiver runs a real time, matched filter detector that continuously seeks matches for tag PN codes. The receivers detect the initial synchronization code by performing an operation similar to that of a GPS receiver in acquisition mode. The GPS acquisition process performs an ambiguity function search, which is a two dimensional search in offset frequency and code phase

Fig 3. Typical receiver setup

space. The animal tracking system simplifies the problem however by eliminating the requirement to search in frequency (Doppler) space, since the expected tag velocities are modest. A full Doppler search can be avoided by keeping the length of the code short, relative to the frequency difference of the reference clocks at the tag and receiver. These slight clock rate differences effectively create small Doppler shifts, and their impact cannot be ignored entirely. The receivers use a GPS disciplined frequency reference which offers sufficient accuracy, however the tags cannot carry a GPS receiver. Inexpensive and accurate TCXO clocks provide the tag reference frequency. Typical tolerances for these clocks are 1-5ppm, corresponding to an offset of 140 to 700Hz at 140MHz. The expected value of the base band signal as a function of carrier phase (Δθ), code phase (Δτ), and frequency offset (Δf_D) are shown in the equation below [11].

$$E\{S\} = \sqrt{C} D e^{j(\Delta\theta + \pi\Delta f_D T_{CO})} \bar{R}(\Delta\tau)\mathrm{sinc}(\pi\Delta f_D T_{CO})$$

The sinc function component can be thought of as an amplitude modulating factor; larger values of Δf_D or T_{CO} reduce the amplitude, which reduces the signal to noise ratio.

While the total tag sequence is 3 msec long, the synchronization and ID codes are processed as separate halves, which yields an integration time, T_{CO}=1.5 msec. If we set $\Delta f_D <= 1/2T_{CO}$, then for the coherent integration time of approximately 1.5 msec, Δf_D can be as large as 333 Hz. Birds carrying the tags will not be moving fast enough to induce an appreciable Doppler shift, therefore this result allows us to avoid the computationally expensive task of searching the ambiguity function in frequency space, provided that the oscillator reference is not worse than 3ppm.

Once the receivers have detected the synchronization code, they then search the next 1.5 msec of signal for all possible ID codes. This search can be done efficiently, since the phase of the code has just been determined. A GPS module in each receive station allows the receivers, which are in different locations, to precisely record the arrival of each synchronization code. This module also provides a disciplined 10MHz clock source for the rest of the system clocks and oscillators.

The matched filter detector offers another advantage in addition to precise time-stamping: it allows us to realize significant signal processing gain. The pseudo-noise code modulation at the transmitter spreads the transmitted energy out over a 2MHz wide band. At the receiver, the matched filter detector de-spreads the signal, effectively reducing the in-band noise, and increasing the signal to noise ratio. This approach allows us to detect signals that are 20 dB below the background noise floor. Conventional receivers, with a human listening to the receiver output, can reliably detect a signal that is 6 to 8 dB below the noise floor. This improved sensitivity enables greater reception range, relative to traditional receivers.

The attractive features of this system, including low power tags, automatic detection, and improved location accuracy, depend on a network of receivers that can listen continuously and in real time for tag transmissions. The real time requirement sets a hard limit on performance which impacts all other design choices. We chose to implement the matched filter detector on a Texas Instruments TMS320C6416 Digital Signal Processor. This processor can run at 1GHz and offers parallel data processing capability, with up to 8GMAC per second when operating on 8 bit data. The incoming RF signal is down-converted to base band I and Q channels and then sampled at 2.8125 MHz. An 11 bit PN sequence at the tag's 1MHz chip rate would occupy 5760 samples, and processing the I and Q channels with a straightforward, time-domain matched filter implementation would require approximately 32 GMAC per second in order to guarantee real time operation. This requirement clearly cannot be met, even by the impressive performance of the 'C6416, which was state of the art when this work began. The options available are to either reduce the bandwidth of the transmitted signal, which reduces the ranging accuracy, reduce the PN sequence length, which reduces the processing gain, or to use a frequency domain algorithm to implement the matched filter. Ordinary time domain correlation is an $O(N^2)$ operation, where N is the number of elements to be cross-correlated. Operation in the frequency domain in contrast is an $O(N\log_{10}N)$ operation, thanks to the remarkable efficiency of the FFT. An early application of this technique to GPS was demonstrated in [12]. We chose this alternative, and used available FFT routines along with the established practice of computing a correlation by conjugate multiplication in the frequency domain, to meet the real time requirements of the system.

The cross correlation function C at a particular offset in samples, k between the signals g(n) and h(n), which are both at least N samples long is

$$C(k) \equiv \sum_{n=0}^{N} g(n)h(k+n)$$

The right side of this equation shows that each output sample (lag) requires N multiplications and additions, leading to high computational costs. An alternative approach uses the following relationship, where * indicates the complex conjugate, IFFT refers to the Inverse Fast Fourier Transform operation, and H(f), G(f) refer to the Fourier transforms of the time series h(n) and g(n).

$$C(k) = IFFT(H(f)^* G(f))$$

This simple equation masks two important caveats: the signals are finite in duration, since they are stored in RAM, and the signals are not periodic. Real-time operation requires that the FFT lengths be as short as possible, since the processing load scales faster than the length of the buffer to be transformed. As an implementation of the Discrete Fourier Transform, the FFT assumes that its input data are periodic in N samples, where N is the length of the input buffer. The correlation technique based on the FFT exhibits a circular behavior, and will "wrap around" data from the end of one buffer onto its beginning for any lag other than 0, where the two buffers are exactly aligned. The solution to this problem is to zero-pad the data buffers at their ends. If +/- k lags are desired, then the data buffers must both be padded with k trailing zeros. The Numerical Recipes book [13] explains this technique in greater detail.

Cross correlation can be thought of as the agreement between two signals, over a window that is N samples long, as they slide past each other. This conceptual image works well if both signals are infinite in length, however end-effects become problematic if the signals are finite. In particular, if both signals are N samples long, and the cross correlation window is also N samples long, any lag other than zero results in samples from one signal being out of range of the other signal: the individual samples at each end of the two signals have no corresponding sample in the other signal. The circular buffer approach, combined with zero padding, resolves this issue, since the end samples are wrapped around onto each other. The wrapped samples are multiplied by zeros, and therefore do not contribute to (or interfere with) the correlation calculation.

Signal-to-buffer alignment is a problem when the signals to be correlated are of equal length. If the start of the incoming PN signal is not exactly aligned with a buffer boundary, the cross-correlation will be lower than the auto-correlation, even when the lag which exactly aligns the signal with the template is chosen. Since the DSP's A/D converter is continuously filling the buffer, the alignment will be arbitrary. In the worst case, the incoming signal is misaligned with the buffer boundary by N/2 samples. In this case, the matched filter detector will still register a maximum at the N/2 lag, but the maximum will be ½ its autocorrelation value, since half of the PN signal is not in the buffer. This reduction in signal strength becomes problematic in low signal to noise situations. A common solution to this problem is to overlap the buffer by 50%, so that the second half of the last buffer becomes the first half of the next buffer that is cross-correlated with the template. This approach requires twice the processing effort of 0% overlap, and it also computes redundant information, since half of the sample data from the previous cross correlation are present in the next correlation.

An alternative approach exploits the time shifting property of the Fourier transform

$$x[n-n_0] \overset{F}{\Leftrightarrow} e^{-j\frac{2\pi}{A}mn_0} X(m)$$

where x is the sampled time series data, n_0 is the number of samples to shift, A is the length of x, and X is the complex spectra of $x[n]$. This property becomes particularly useful when the time shift is $A/2$, since the complex exponential reduces to the sequence [1, -1, 1, …]. This sequence can be stored in memory, rather than computed at run time. Multiplying $X(m)$ by this simple sequence yields the same spectra as would shifting $x(n)$ by n_0 and recomputing the Fourier transform (note that these shifts are circular shifts within the buffer). In fact, no actual multiplications need take place at all, since this is merely a sign change on every other data entry.

The implementation of this fast correlation algorithm attended to two issues: avoiding artifacts from the circular nature of correlation using FFTs, and minimization of the number of forward and inverse FFTs that must be computed. The data buffers are shown in figure 4; each buffer is identified by a roman numeral. The figure depicts the data in the time domain for clarity, though they are processed in the frequency domain. To avoid artifacts from circular correlation, the length N/2 code template is zero padded to length N, and its FFT is precomputed and stored as a constant complex vector (buffer V). This length N/2 template is correlated against length N segments of real data (buffer IV) using the fast algorithm, but only the first N/2 lag terms are used. The continuous input data stream (buffer I) is segmented into length N/2 segments, referred to as frames. Each current frame (B_0) is processed as it becomes available from the A/D. A copy of the current N/2 length frame is made and padded with N/2 zeroes (buffer III). The copy of the previous frame is preserved (buffer II). To form the DFT of the length N concatenation of these two input segments, the DFT of the first segment is added to a modified version of the second DFT. The modification is simple: the sign of the even numbered elements of the vector is reversed. This sum provides the FFT of the length N sequence for termwise conjugate multiplication in the fast correlation algorithm (buffer IV). Note that in processing the next frame, the DFT of the second segment is used again, unmodified, to form the first half of the next length N segment for processing. A fresh set of N/2 samples is read in, zero padded, processed by FFT, and modified by sign reversals before adding to the FFT of the second block. Thus, each set of N/2 correlation lags requires one input FFT, N/2 sign changes, N complex multiplications, and one inverse FFT. This block processing algorithm provides cross correlation output that is free from gaps or artifacts.

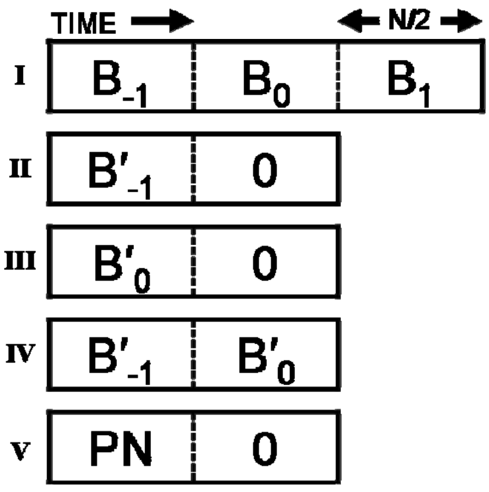

Fig 4. Buffers used in fast FFT-based correlation engine

D. Networking, Data Processing, Storage

In addition to the DSP and GPS module, each receiver contains an embedded Linux Single Board Computer (SBC) which is used to capture DSP detection strings through an RS-232 interface. Multi-threaded Python code parses the serial input and attempts to submit the data via an Ethernet link to a remote server. The Ethernet port on the SBC is connected to a GSM modem with a crossover cable. This enables wireless Ethernet connectivity within the Cingular/AT&T service area. Data are stored locally in non-volatile compact flash memory using an SQLite database. If the connection to the server is interrupted, stored records are retransmitted when wireless communication with the server is reestablished. Data records are bundled and zlib compressed for submission over a TCP stream to the server. The Twisted Python networking library was used to write a TCP server process and protocol that is capable of handling simultaneous data streams from many receive stations. The server process receives batches of compressed data events and submits them into the central PostgreSQL database for localization and analysis.

E. Localization and error

As mentioned earlier, the system can determine a position fix with 3 arrival time measurements. This implies that the position fixes are established in a 2 dimensional plane, and are not full, 3 dimensional fixes. Finding the unknown position of a transmitter in 3 dimensions involves solving a set of equations for tag position (x,y,z) as well as transmit time. Thus, these four unknowns require four or more equations for a solution. In our system, the position is assumed to lie in the xy plane and therefore only 3 equations are necessary. This allows the system to achieve position fixes with only 3 towers, though additional towers make the fixes more robust, and increase the accuracy. The assumption of a planar solution will naturally lead to positioning errors when the source (tag) is not exactly in-plane. The receive tower geometry is also unlikely to be exactly planar. These errors are small when the system is used in relatively flat areas, though they cannot be neglected in hilly applications. In these cases, the full 3D solution must be found, which requires four or more simultaneous tag detections.

When tag transmissions are detected by the receivers, the arrival times are routed through an auxiliary data network to a central server. There, time of arrival information for the received events is passed into a localization algorithm implemented in the MATLAB computing environment. This algorithm estimates transmitter location by performing an iterative spatial search over the animal study area that minimizes the sum squared error between the actual TOA vector and that of candidate locations in the search space. The search uses a weighted stochastic method that balances convergence speed and robustness to eccentricities in the data. A full description and discussion of this method, with comparison to others, is the subject of a forthcoming paper.

As is often the case in arrival-time-based positioning systems, system geometry has a significant impact on our system's accuracy. Though we do not use the positioning algorithms typically employed by GPS receivers, the concept of dilution of precision (PDOP) still provides a useful assessment of the likely sensitivity to errors in the system due strictly to the geometry. A cursory review of the derivation of PDOP is provided next. Further details may be obtained from [11].

GPS receivers solve an over determined estimation problem for position and time, based on pseudorange measurements. The term pseudorange describes the range from a satellite to a receiver, and is measured by the receiver. In the GPS system, the measured pseudorange (based on known transmit time and measured arrival time) to each k satellite is

$$\rho_c^{(k)} = \left\| \mathbf{x}^{(k)} - \mathbf{x} \right\| + b + \varepsilon^{(k)}$$

where $\mathbf{x}^{(k)}$ and $\mathbf{x}$ are the three dimensional position vectors of the kth satellite and the receiver, respectively, b is the pseudorange error (in meters) due to receiver clock offset from correct time, and and $\varepsilon^{(k)}$ is an error term for each satellite. To be clear, $\mathbf{x}$ and b are the variables to be solved for. The GPS system also approximates the pseudorange to each receiver as

$$\rho_0^{(k)} = \left\| \mathbf{x}^{(k)} - \mathbf{x}_0 \right\| + b_0$$

where the new term $\mathbf{x_0}$ is current (possibly wrong) guess of the receiver position, and b_0 is the current guess for the receiver clock offset, again expressed in meters. The new δ terms

$$\delta\rho^{(k)} = \rho_c^{(k)} - \rho_0^{(k)}, \delta\mathbf{x} = \mathbf{x} - \mathbf{x}_0, \delta b = b - b_0$$

are the differences as shown, and are used as error terms. It can be shown [11] that

$$\delta\rho^{(k)} \approx -\frac{(\mathbf{x}^{(k)} - \mathbf{x}_0)}{\left\| \mathbf{x}^{(k)} - \mathbf{x}_0 \right\|} \cdot \delta\mathbf{x} + \delta b + \varepsilon^{(k)}$$

where a Taylor series approximation of a vector norm has been used. The same equation may be written in matrix notation as

$$\delta\boldsymbol{\rho} = \mathbf{G}\begin{bmatrix} \delta\mathbf{x} \\ \delta b \end{bmatrix} + \varepsilon^{(k)}$$

GPS then proceeds to solve for

$$\begin{bmatrix} \delta\mathbf{x} \\ \delta b \end{bmatrix} = \mathbf{G}^{-1}\delta\boldsymbol{\rho}$$

however we have employed a different method to solve these equations, as mentioned earlier. The matrix **G** however, which is based solely on the positions of the satellites and the receiver, can also be used for our system, where the role of receiver and transmitter is reversed. The 4 X 4 matrix **H**, defined as

$$\mathbf{H} = \left(\mathbf{G}^T\mathbf{G}\right)^{-1}$$

can be used to determine the dilution of precision (PDOP) as

$$(\mathrm{PDOP}) = \sqrt{\mathbf{H}_{11} + \mathbf{H}_{22} + \mathbf{H}_{33}}$$

In our system, the only difference is that **H** will be 3 X 3, since we are only computing position in x and y. One standard deviation of position error is then

$$\text{Position Error} = \sigma \cdot \text{PDOP}$$

where σ is the ranging error, a term dependent on several system parameters, including chip rate and signal to noise ratio, but not dependent on system geometry. Therefore we can use PDOP to assess the likely quality of the locations at different positions in the covered area, and make system design tradeoffs.

Hyperbolic intersection [14] methods have been widely used in the animal tracking community to localize animals based on time of arrival measurements, particularly in the bioacoustics field. This graphical, and therefore potentially more intuitive method can also be used to demonstrate the error sensitivity due to system geometry. The locus of potential source locations for a unique time difference of arrival measurement at two physically distinct receivers falls along a hyperbola in 2D. A different pair of receivers will yield a different hyperbola. The intersection of these curves yields the position estimate. Though our animal tracking system is not using the traditional crossed hyperbola approach to estimate position from time delays, this graphical approach can be helpful in building intuition for the impacts of system geometry on performance. A simple example is shown in Fig 5. Receiver positions are shown as filled circles, while the two source locations, labeled 1 and 2, are shown as un-filled circles. The lines of

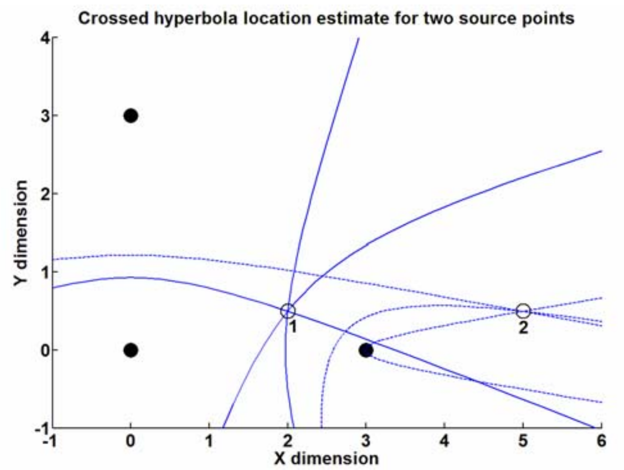

Fig 5. Method of crossed hyperbolae for determining position

hyperbolae which correspond to each source location have been shown. This method uses the intersection of the hyperbolae to estimate the source location. Notice that as a source point moves "outside" of the receiver array, the lines of hyperbolae intersect the transmit location at increasingly acute angles. While the true reason for this is that the geometry of the system causes the time difference between pairs of receivers to be at or near the maximum, and the incremental change in time delay due to a slight displacement of the source to be at a minimum, the impact in the plot is to reduce the angles of intersection. Small timing errors will perturb each hyperbola slightly, which translates into a relatively larger position error, since the point at which the hyperbolae cross becomes increasingly sensitive as the angles of intersection become smaller. While our approach, based on the iterative spatial search, seems to outperform the simple crossed hyperbolae algorithm, both algorithms, like the approach used in GPS, are still based on arrival time differences, and are therefore subject to the same physical sources of error. As a result, no matter which error estimation tool we use to assess the array performance, the performance of the system is far better if operation can be restricted to source locations "inside" the array. In practical terms, this means that a receiver system intended to cover a large area must use a tiled approach, with receivers arrayed at regular intervals. The distance between receivers is determined by site specific parameters, including topography and foliage. The present system is capable of operation in deciduous forest with tag to tower ranges up to 7 km. If installed in open areas, or on hilltops, we expect the range to increase. Efforts to improve the range through increasing processing gain and reducing interference are ongoing.

III. SYSTEM PERFORMANCE

Our first tests were designed to evaluate range of detection under realistic conditions. A single transmitter, configured to transmit once per second, was moved successively farther away from the receiver until only half of the transmissions in a five minute interval were detected. These tests were carried out over flat terrain and required the signal to pass through mixed deciduous forest during the summer months. We found, unsurprisingly, that the range was most strongly dependant on the tag's height above ground, the antenna orientation, and the topography between the transmitter and the receiver. Though the maximum range of detection varied widely, we used an average measure of 7 km for receiver network planning. This range is strongly site-dependent, and will need to be reevaluated for new sites.

In addition to range testing, we also carried out controlled tests to measure the minimum required signal power for detection. A lab test version of the transmitter, configured for 50 ohm output impedance, was connected to a variable attenuator. The attenuator and tag were then connected to the receiver via a 10 meter length of LMR-400 cable, and the signal attenuation was increased until the receiver no longer detected the signal (the detection criteria was the same as in the range test). The tag's output power, which is adjustable, was set to 12dBm (this is the total power in the 2MHz wide main lobe). The maximum attenuation that still allowed reception events was -136 dB, yielding a minimum sensitivity of -124dBm for our current receivers. Though this number could, in principle, be used to formulate a link budget, uncertainty in the tag antenna's gain prevents more detailed estimates of range. As mentioned earlier, the tag's present antenna uses an electrically short whip. The expectation for this antenna is that the bird carrying the transmitter forms a reasonably good ground plane. This assumption proves to be untrue in many cases, leading to poor radiation performance and output stage impedance mismatch. This antenna problem remains unresolved, appears to contribute to widely varying range of detection performance, and is therefore a priority for our future work.

We installed a four receiver system in Ithaca, New York to test and validate our localization approach. Table I shows the results of a five minute long static test conducted near the center of the array. The tag was configured to transmit once per second, therefore there were 300 computed positions during this interval. A test position was chosen near the center of the array, since it should yield system geometry with low dilution of precision. A consumer grade GPS was used to establish the location of the tag and the locations that the GPS and the Auto-locating system reported are given. Unfortunately, the GPS did not enable explicit averaging, and its displayed estimate of error was approximately 5 meters during the test. The position estimates of the GPS and the tracking system agree closely in Easting, but differ by 19 meters in Northing. There are several

TABLE I
TEST RESULTS FOR FIXED TRANSMITTER

	GPS	Auto-Locating System	
	Avg Pos (km)	Mean Pos (km)	STD (km)
Easting	381.426	381.425	0.0091
Northing	4700.084	4700.065	0.0089

possible explanations for this discrepancy, including noisy GPS position estimates, higher DOP in y than in x, noisy receiver location estimates, and slightly out of plane receiver and transmitter locations. Figure 6 shows a histogram of the estimated Easting (x) coordinates from this test. The histogram for Northing (y) (not shown) indicates a comparable spread. The STD, which is about 9 meters in this test, indicates that the system is capable of positioning an animal inside a circle which is roughly 10 meters in radius. This is excellent resolution for

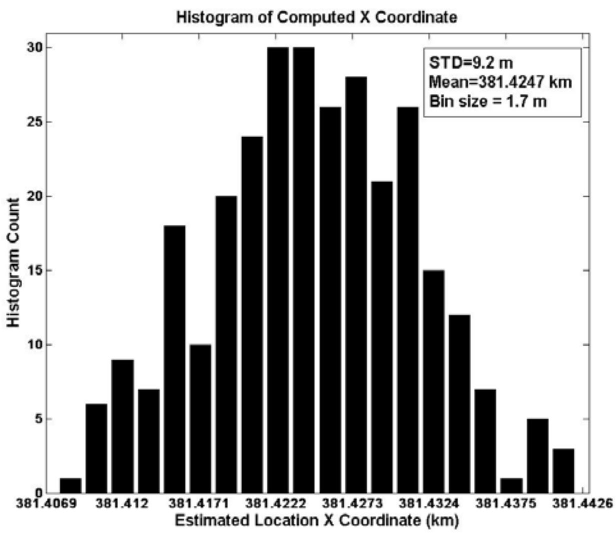

Fig. 6 Histogram of computed Easting (x) Coordinate

Fig. 7 American Crow with prototype DSSS transmitter

Fig 8. Track of Crow WJ

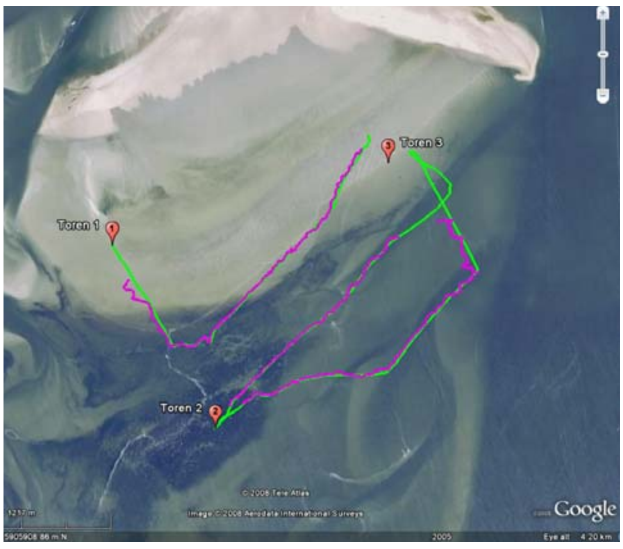

Fig 9. Boat track in Wadden Sea comparing Automatic System and GPS

mapping territories and characterizing patterns of movement, but finer resolution or onboard sensors would be required to begin capturing more detailed aspects of behavior. Kalman filtering techniques may improve the resolution.

We captured six American Crows and affixed the tags using a backpack style harness (Fig. 7). The tags were configured to transmit once per second, in order to ensure sufficient data for post-processed analysis. The data for the arrival times of each tag were aggregated in a database and then processed with our position estimation software. The computed locations were overlaid on Google maps-based images as shown in Fig. 8. The positions of the receivers are indicated with orange circles, and the track of the crow is shown with a red dashed line.

We also conducted tests in the Dutch Wadden Sea. Figure 9 shows the track of a small boat, which carried a transmitter as well as a consumer GPS receiver. The track generated by the system is magenta, while the GPS track is green. The towers are represented with red pushpins, and are roughly 2 km apart. The two tracks agree closely, with minimal offset bias and noise. The system did not perform as well when the transmitter was extremely close to any of the towers; the generated track showed significant "noise", centered on the actual transmit position. For clarity, the data from these locations has been clipped out of the image. This situation may be similar to observed "close-in" errors in the pseudolite community, and we are investigating potential solutions. This issue is not likely to be problematic in practice, since any field application will employ a tiled receiver arrangement, which guarantees that, except on the array boundaries, more distant, and therefore better positioned receivers will also detect a transmission.

IV. Conclusions

Conventional radio tracking techniques are not suitable for monitoring large numbers of animals over long time scales. They are labor intensive, yield relatively few position fixes per unit time, and use their batteries inefficiently. Existing automatic location tags based on GPS and ARGOS are too heavy for many birds. We have demonstrated a system for tracking large numbers of animals accurately over medium ranges. The transmitters are very small, making them suitable for most birds, and can last several years.

Acknowledgment

We would like to thank Jim Omura for technical advice and the Gordon and Betty Moore Foundation for supporting this work. Kevin McGowan and Anne Clark trapped and tagged the crows in this study. Theunis Piersma and NIOZ supported the Wadden Sea tests. We are also indebted to the Hartmann, Brenna, and Todhunter-Fubini families who graciously allowed us to place and maintain our receivers on their land.

References

[1] C.D. LeMunyan, W. White, E. Nybert, "Design of a miniature radio transmitter for use in animal studies," *J. Wildl. Mgmt*., vol. 23(1), pp. 107-110, 1959.

[2] W.W. Cochran, R.D. Lord, "A radio tracking system for wild animals," *J. Wildl. Mgmt*., vol. 27(1), pp. 9-24, 1963.

[3] J.E. Lee, "Accessing Accuracy of a Radiotelemetry System for Estimating Animal Locations," *J. Wildlife Manage*. vol. 49, pp. 658-663, July. 1985.

[4] W.W. Cochran, D.W. Warner, J.R. Tester, V.B. Kuechle. "Automatic Radio-Tracking System for Monitoring Animal Movements," *BioScience*. Feb, 1965

[5] M. Green, T. Piersma, J. Jukema, P. De Goeij, B. Spaans, J. Van Gils. "Radio-telemetry observations of the first 650 km of the migration of Bar-tailed Godwits *Limosa lapponica* from the Wadden Sea to the Russian Arctic," *Ardea*, vol. 90(1), pp. 71-80, 2002.

[6] Cochran, W.W., G. Swenson, Jr., L. Pater. (2002) "Radio Direction-Finding for Wildlife Research" [Online]. Available:http://userweb.springnet1.com/sparrow/ Direction-finding.html (2008) Automated Radio Tracking System (ARTS). [Online]. Available: http://www.princeton.edu/%7Ewikelski/research/index.htm

[7] K.J. Gaston, T.M. Blackburn. "The frequency distribution of bird body weights: aquatic and terrestrial species". *Ibis* 137:237-240. 1995.

[8] S. Cobb. "GPS Pseudolites: Theory, Design and Applications", Ph.D. Thesis, Stanford University, 1997.

[9] J. Strada, D. Henderson. "Navstar Field Test Results". In Global Positioning System vol 1, Institute of Navigation, Washington, D.C., 1980.

[10] R. Gold, "Optimal Binary Sequences for Spread Spectrum Multiplexing," *IEEE Trans. On Information Theory,* vol. IT-13, pp. 619-621, Oct. 1967.

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

[12] DJR. Van Nee, AJRM. Coenen. "New Fast GPS Code-Acquisition Technique Using FFT" IEEE Electronics Letters, Vol. 27 No. 2, 17th January 1991.

[13] W. Press, S. Teukolsky, W. Vetterling, B. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed., Cambridge University Press, 1992.

[14] B. Fang. "Simple Solutions for Hyperbolic and Related Fixes" *IEEE Transactions on Aerospace and Electronic Systems*, vol. 26, no. 5, Sept. 1990.