Untying the knot

Bijleveld, Allert Imre

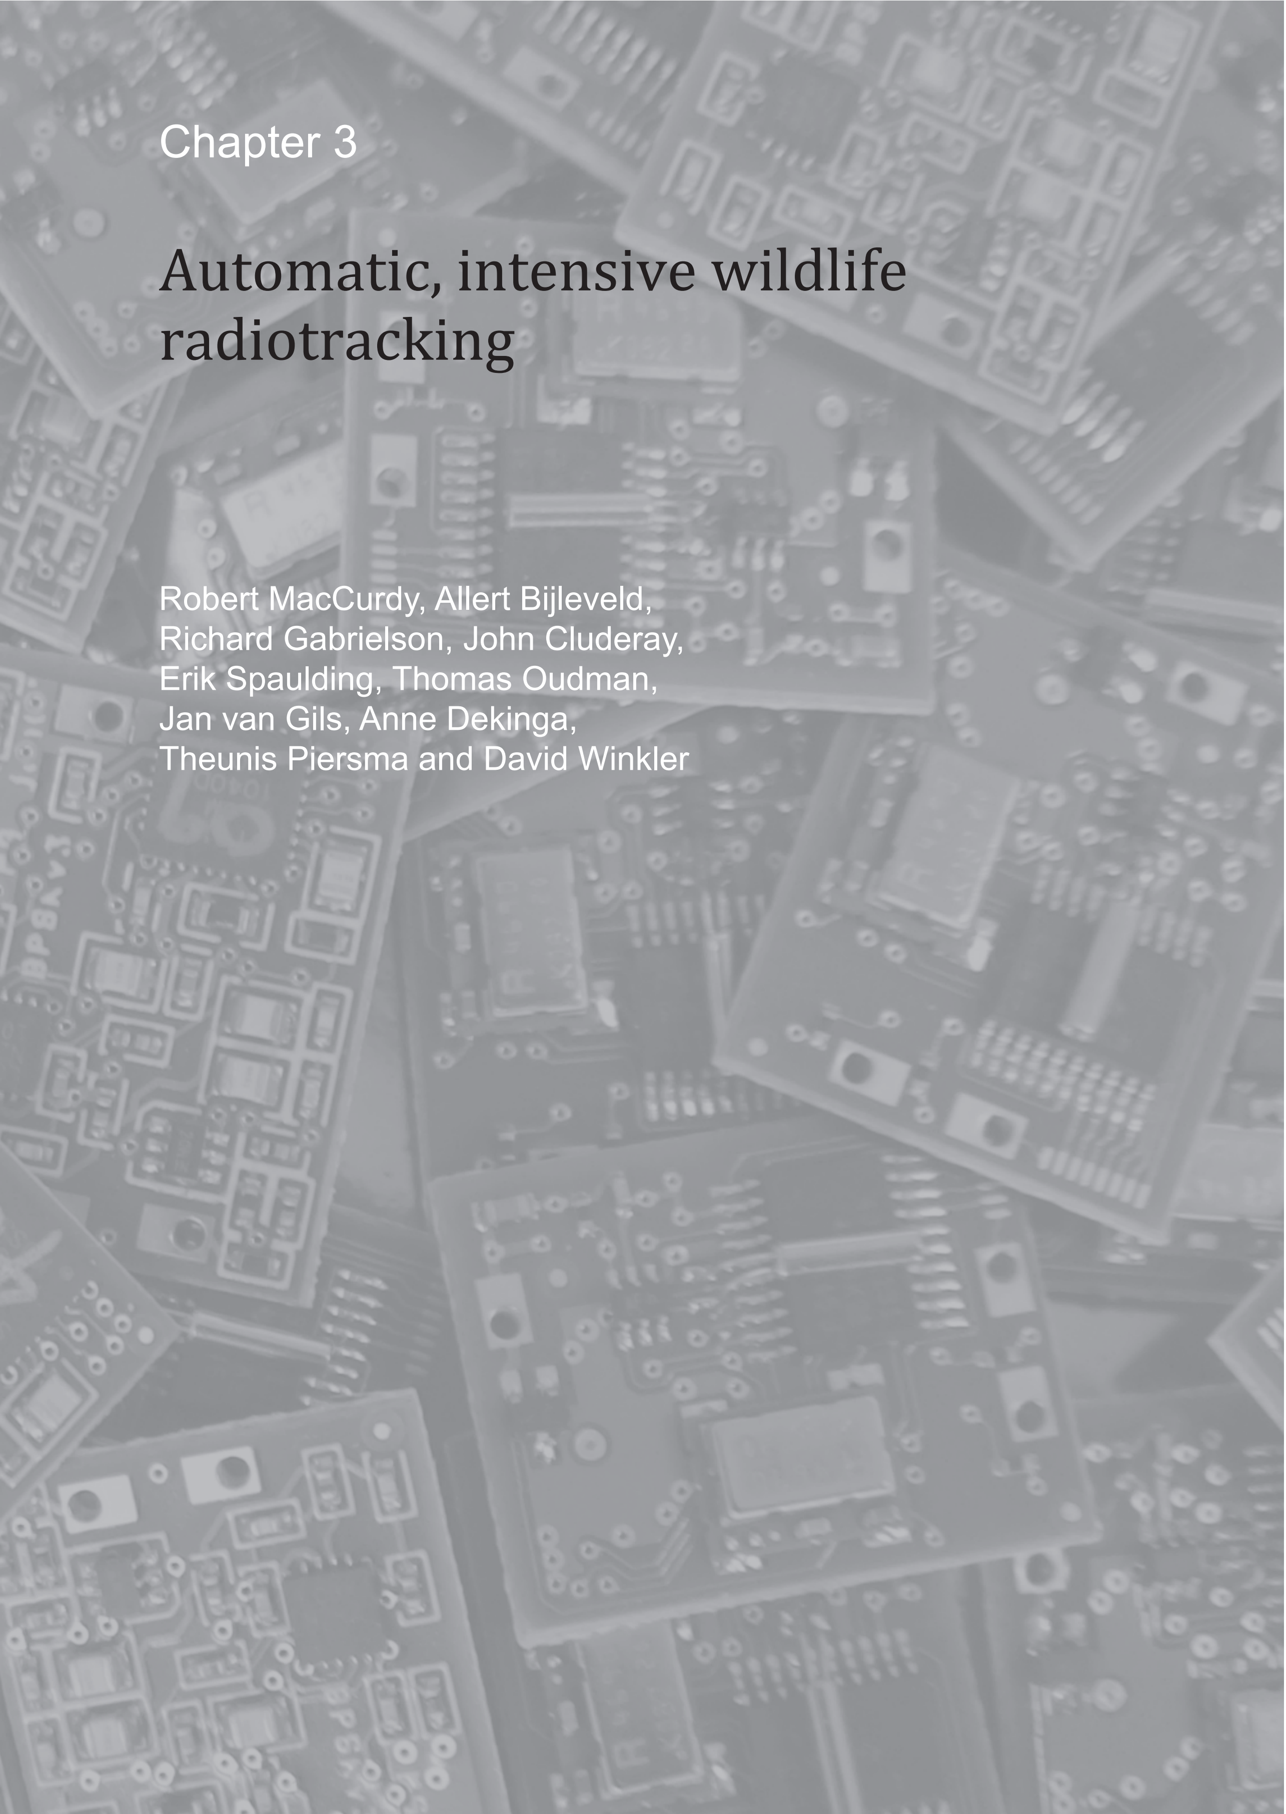

Chapter 3

Automatic, intensive wildlife radiotracking

Robert MacCurdy, Allert Bijleveld,
Richard Gabrielson, John Cluderay,
Erik Spaulding, Thomas Oudman,
Jan van Gils, Anne Dekinga,
Theunis Piersma and David Winkler

SUMMARY Recent advances in tracking technology, in part enabled by the explosion in personal wireless connectivity, have begun to give wildlife scientists the scalable tools required to monitor large numbers of individuals. Unfortunately, many of these new tools are inapplicable to many species due to mass, cost and energy constraints, leaving gaps in our understanding. Here we present a new technique, capable of automatically gathering position data with high spatiotemporal resolution for large numbers of animals over long timescales, using very small transmitters. Relative to current methods this system offers researchers unprecedented amounts of data, can be broadly applied to species that were previously too small for automated tracking systems, and reduces tracking costs. We describe the challenges encountered when tracking wildlife with existing technologies, our solution as implemented, and discuss application examples.

INTRODUCTION

Movement is fundamental to all living organisms and its study is used directly and as a proxy to quantify diverse parameters across spatial and temporal scales. The Movement Ecology field seeks to provide a unifying framework, including methods, for the long-established but often disparate practices of investigators studying organismal movements (Sugden and Pennisi 2006, Nathan et al. 2008). Researchers studying the dispersal of maple leaves, the advance of invasive insects, or the spatial resource utilization of foraging animals have traditionally developed or purchased tools specifically for their application area, with little opportunity to share these tools with groups working on different ecological systems. This practice is changing; the National Science Foundation's National Ecological Observation Network (NEON), which represents a $400 million investment over five years, is one example. NEON aims to develop ecological sensing infrastructure at 20 locations around the United States (Keller et al. 2008, Pennisi 2010), and to allow researchers to share field resources and collected data. These locations will persist for 30 years, provide year-round power, internet connectivity, and host a variety of automated and staffed sampling tools. Though its current development plan does not incorporate automated movement monitoring tools, the NEON sites, and many others around the world, present movement ecologists with an appealing opportunity: when instrumented with automated tracking tools, pervasive, persistent, coordinated and automated data collection sites have the potential to dramatically enhance our understanding of organismal movements through space and time.

We have developed a Real-Time Locating System (RTLS), based on a Time of Arrival (TOA) approach, capable of monitoring the positions of thousands of wildlife transmitters (tags) in near-real-time. This terrestrial tracking system targets regional coverage, and complements global systems like GPS by dramatically expanding the number of species that can be tracked, increasing tag lifetimes, raising location update rates, and reducing per-animal tracking costs. This method increases the number of position estimates that can be obtained from small tags by several orders of magnitude, relative to existing techniques (see Figure 3.1). In contrast to approaches that utilize existing data networks (satellite, mobile phone), by employing local point-to-point wireless connectivity, this method offers real-time position updates without recurring data costs. Since most of the system's cost is in the fixed receiver network, the incremental cost of adding additional tagged animals to the study area is extremely low, which enables large sample sizes that would be impractical or impossible via existing methods, and opens up the potential for shared tracking infrastructure.

TAG DESIGN CHALLENGES

Wildlife tracking systems that are capable of providing position information are a widely used tool; however, their application is limited by cost and mass to a relatively small number of species. To illustrate why a new tracking technology is required, the character-

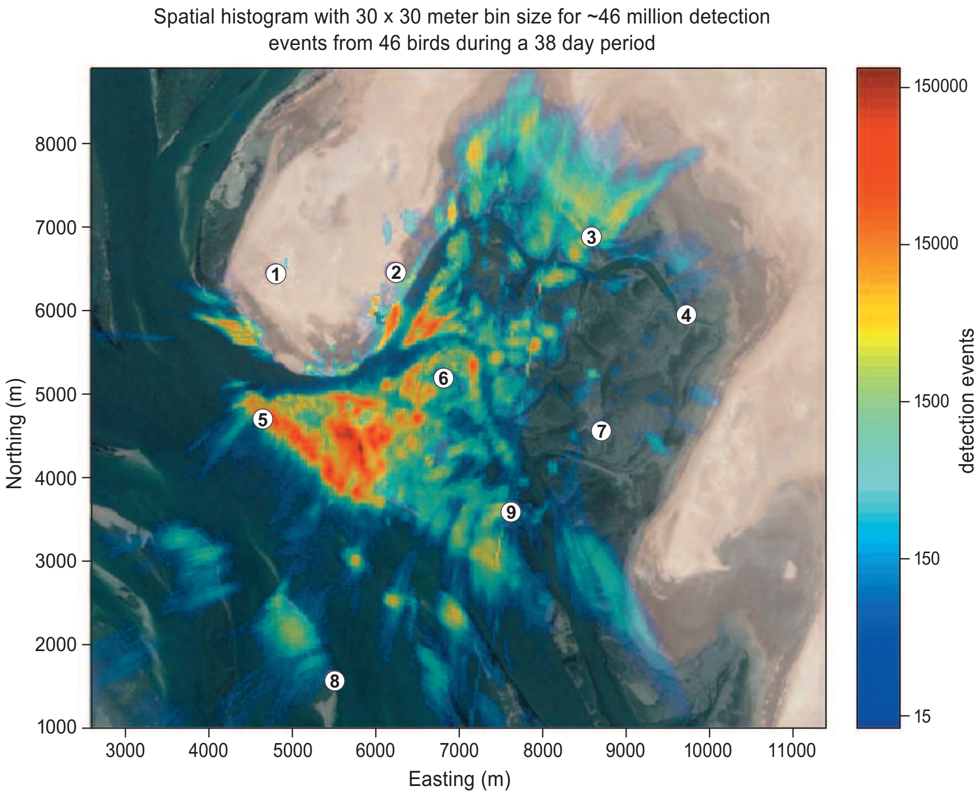

Figure 3.1 Baie d'Aouatif on the Banc d'Arguin, Mauritania with a spatial histogram showing the number of detection events in each 30x30 m cell. 46 birds were tracked over a 38 day study period producing more than 46 million position estimates. The TOA tags were configured to transmit once per second. The numbers represent the nine receiver stations.

istics of existing systems are summarized in Table 3.1. Though numerous tag sensing modalities exist, here we consider only those that provide position data. Of the parameters listed in Table 3.1, the most important is tag mass, as it determines whether a tracking methodology is suitable for any particular species. The precise amount of tag mass that a given animal can bear without adverse impacts is still unknown. Though efforts have been made to estimate the allowable load based on species-specific parameters (Caccamise and Hedin 1985), it is common practice to limit tag mass to 3–5 percent of an animal's body mass (Cochran 1972, Naef-Daenzer et al. 2001). This heuristic has a profound impact on tag design. As illustrated in Figure 3.2, tags must be lightweight if they are to be compatible with most flying vertebrates. For example, a 0.85 gram tag is light enough to be used with half of the animals included in Figure 3.2 when a 3% loading rule is enforced; for comparison the lightest Argos (Fancy et al. 1988) (satellite tracking) tag available, a 5 gram model, can be applied to fewer than 20 percent of the species presented in Figure 3.2. The increase in the number of species that can be tagged increases most rapidly for tag

Table 3.1 Comparison of different tracking technologies. All numbers represent the most favourable values currently available. NOTES:

A Tag range can be limited either by the position sensing mechanism or by the data offload system (if present). The minimum of these two is shown.
B Location Cost - total tag energy cost of each position fix. For tags that only transmit, this entry is zero. For GPS this metric is strongly dependent on the position update rate. Faster update rates use less energy per fix, but cannot use low power sleep modes, yielding greater overall energy usage than slower update rates. These data are taken from published research as well as manufacturer's data sheets; efforts have been made to provide a range from best to worst case.
C Data Transport Cost - total tag energy cost of transmitting the data or signal for each position fix. Total cost is the sum of the Location Cost and the Data Transport Cost.
D Lifetime w/ 1g battery - number of days that a tag could operate when acquiring 1 position fix per minute, using a 1 gram (35mAh) battery. For all tags that support it, this calculation includes the transport cost of sending the data back to the user, rather than storing locally on the tag, since recapture is often infeasible. The VHF calculations were done with a pulse interval of 3 seconds, which is typical for very small "Beeper" VHF tags.
E Some Argos and GPS tags are available with solar cells, which can extend their lifetime indefinitely, though the number of fixes per day is strongly linked to insolation and the strength of the GPS signal; Bouten et al. (2013) report results varying from 15 to nearly 7,000 fixes/day.
F The GPS (research) row is an estimate, based on the best recent reported results for GPS baseband and RF frontend modules.
G Conventional VHF tags usually send a simple presence/absence signal, but do not send modulated data.

	Tag Cost (USD)	Minimum Mass (g)	Range $(km)_A$	Location Error (m)	Data Back haul	Location Cost $(J)_B$	Data Transport Cost $(J)_C$	Lifetime w/ 1g $battery_D$
GPS[1]	50 - 1000	1 - 2.5	Global	1 - 50	N	0.022 - 2	0	0.58 $days_E$
GPS & Argos[2]	1500-4000	17	Global	1 - 50	Y	0.022 - 2	0.38	0.32_E
GPS & local wireless[3]	n/a	6 - 12	0.3 - 8.5	1 - 50	Y	3.3 - 145	n/a	n/a_E
Snapshot GPS[4]	3400 - 5000	39	Global	1 - 100	N	2.1	0	0.12 $days_E$
GPS(research)[5] $_F$	n/a	n/a	Global	1 - 50	n/a	0.26	n/a	1.0
Argos[6]	1300 - 4000	5	Global	250 - 1500	Y	0	0.18	1.45 $days_E$
Solar Geolocation[7]	100	0.5	Global	50×10^3 - 2×10^5	N	6.9×10^{-5}	0	3804 days
VHF[8]	150	0.16	1 - 5	100 - 1000	Y_G	0	1.0×10^{-4}	130 days
TOA RTLS	150	0.5	5 - 10	10 - 50	Y	0	5.0×10^{-4}	520 days

1 Technosmart tracking systems for animals, http://www.technosmart.eu; Lotek wireless fish & wildlife monitoring, http://www.lotek.com; Telemetry solutions, http://www.telemetrysolutions.com; Microwave telemetry, http://www.microwavetelemetry.com; Bridge et al. (2011); Edwards et al. (2013)
2 Microwave telemetry, http://www.microwavetelemetry.com; North star science and technology, http://www.northstarst.com; Fancy et al. (1988); Bridge et al. (2011)
3 Bouten et al. (2013)
4 Liu et al. (2012)
5 Tang et al. (2012); Heiberg et al. (2011)
6 Fancy et al. (1988); Priede and French (1991); Hooijmeijer et al. (2014); Battley et al. (2012); Gill et al. (2009)
7 Hill (1994); Afanasyev (2004); Bridge et al. (2013)
8 LeMunyan et al. (1959); Cochran and Lord Jr (1963)
9 MacCurdy et al. (2009, 2012); Piersma et al. (2014); Savaglio et al. (1997); *Lemnell et al. (1983)*

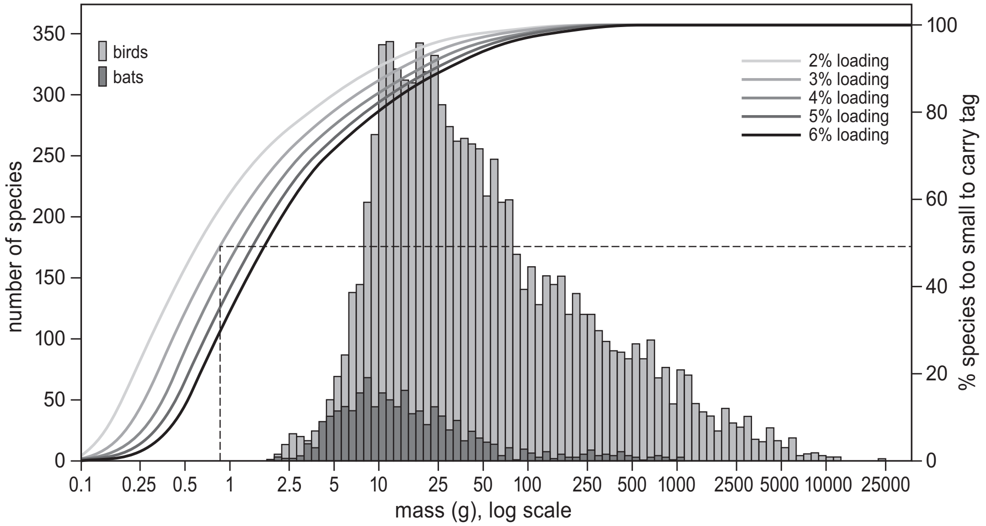

Figure 3.2 Bird and Bat mass distributions; 9278 bird species and 1067 bat species are included. Shown on the right y-axis is the percentage of birds and bats that are too small to carry a tag with a particular mass, assuming a certain maximum allowed loading. For example, the dashed lines show that 50% of birds and bats are too small to carry a 0.8 g tag, assuming 3% loading. Data are from (Smith et al. 2003, Dunning 2008).

weights between 0.5 g and 3 g; this observation motivates the design of a new class of ultra-lightweight tags.

Although we emphasize tag mass here, recent work has shown that aerodynamic drag due to tag cross-section can be significant for small birds, and should not be ignored (Bowlin et al. 2010). Nevertheless, lighter tags will generally be smaller and exhibit lower drag, so the primary focus must be to reduce the size and mass of tag components. As the largest and heaviest components, batteries dominate the mass budget of most tag designs. The energy storage capacity of any particular battery type is proportional to its mass, and steady improvements in mass-and volumetric-specific energy have enabled ever smaller batteries; however, unlike the exponential growth of transistor counts and attendant improvements in power consumption that integrated circuits have exhibited, the battery development curve over the past 20 years has shown only linear improvements (2.5x improvement) (Oudenhoven et al. 2012). Recent efforts (Chen et al. 2014) have yielded micro batteries with more than double the mass-specific energy density of comparable commercial offerings; however, if the past trend is predictive, commercial battery technology is likely to improve these metrics by only 125% over the next decade. As the three right columns in Table 3.1 reveal, this improvement will not be sufficient to allow more energy-intensive techniques like GPS to become incorporated into small tags. An improvement of more than three orders of magnitude in either circuit efficiency or energy storage will be necessary if tags that use GPS or Argos are to become both very lightweight (<1g) and long lived (100 days or more of operation).

To accommodate the limitations of current battery technology, tag designers can reduce operational lifetime, reduce the energy-intensity of the position sensing modality, or capture external energy to replenish the battery. This latter approach, utilizing photovoltaic cells, is now commonly exploited by tag designers with good results. However, the amount of incident power available to the solar array, which must be small enough to accommodate tag size and mass constraints, is insufficient to directly supply higher-power position determination schemes like GPS and Argos. For example, one of the most efficient solar cells currently available (Spectrolab tasc solar cell) has a maximum rated output of 0.027 W/cm^2, occupies 2.277 cm^2 and weighs 234 milligrams, while one of the lowest-power commercially-available GPS modules (SiRFstarIV) consumes 0.077W during acquisition. A tag using these components would require 3 cm^2 of solar cell area, weighing 300 milligrams. This is an ideal-case estimate, and the amount of solar power available is usually dramatically reduced by habitat characteristics, weather, time of day, season and component degradation, which would require a solar array at least an order of magnitude larger than this estimate. To circumvent these limitations, tags employ energy storage elements (batteries or capacitors) to accumulate solar energy over time and then rapidly discharge this energy in the position sensing circuitry. Rather than considering the power balance, tag designs must consider the daily energy balance. Solar cells can yield GPS and Argos tags that are capable of very long deployed lifetimes, but since the tag energy demands are relatively high and the external energy supply varies, these tags exhibit high variability in the number of position fixes per unit time (Bouten et al. 2013). Future GPS designs might one day improve this situation through lower power operation, though it is instructive to look at recent trends. A state of the art research (not commercialized) GPS-receiver published in 2000 (Namgoong et al. 2000) consumed 21.3 mW during the continuous tracking stage. Over a decade later, the lowest power GPS receiver demonstrated in a research setting consumed 8.7 mW (Cheng et al. 2009, Tang et al. 2012) during continuous tracking. For comparison, one of the lowest-power commercially available GPS receivers (SiRFstarIV) consumes 66 mW during continuous tracking. Decoding GPS signals is an inherently compute-intensive operation, and while the energy-intensity of GPS receivers continues to decline, it seems unlikely that energy reductions of 2 to 3 orders of magnitude will occur in the near future.

In light of these current limitations, the choice that tag users face is to trade-off tag functionality for lifetime, since tag mass is a fixed constraint for any particular species (see Table 3.1). Here, functionality could mean various things depending on the tag, including: number of position fixes per day (GPS & Argos), number of light measurements per day (Geolocators), and number of pulses per second (VHF ”beeper” tags). Reducing the number of data points gathered per day does increase the tag endurance, but at the cost of lower temporal resolution, an issue that precludes this strategy for many studies. Though they offer global coverage and high resolution, GPS-and Argos-based tags suffer from high energy consumption, requiring larger and heavier batteries. GPS receivers yield position fixes that are local to the tag, requiring additional tag energy to telemeter the position data. Geo-location tags are extremely energy-efficient and lightweight, but provide very

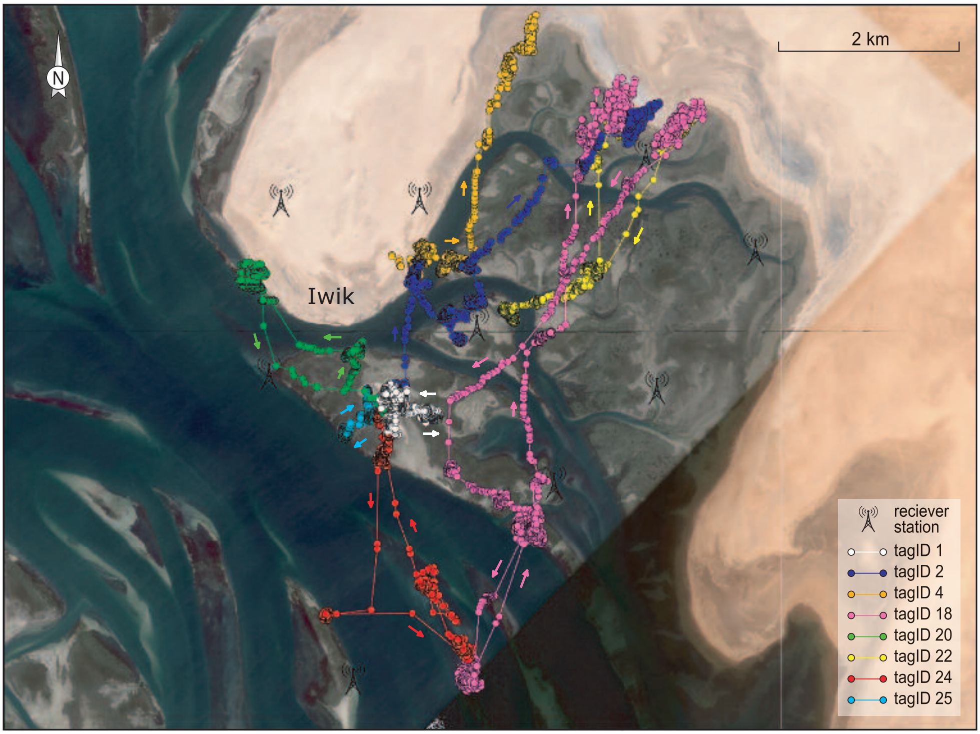

Figure 3.3 The movements of eight Red Knots (*Calidris canutus canutus*) in Baie d'Aouatif on the Banc d'Arguin, Mauritania, during a single low water period (29 January 14:52 h to 30 January 03:07 h) that show the range of individual itineraries. Connected dots indicate the measured positions and the arrows indicate the directions of movement in the course of the tide. Locations that are well separated probably indicate flight paths. This image was adapted from Piersma et al. (2014).

coarse resolution and rarely offer the capability to telemeter data. Conventional VHF tag tracking systems do not provide long tag lifetimes or offer precise position estimates. By minimizing the energy consumed per position estimate, our TOA tracking solution is able to provide remote (tag recapture is not required) position updates at high, consistent rates over long study periods using lighter tags than any other method, allowing a dramatic increase in the number of addressable species (Figure 3.3).

APPLICATIONS

Climate change, and human encroachment on critical habitat place ever-increasing pressure on wildlife, yet the evidence of specific impacts often arrives too late (Boere and Piersma 2012, Pimm et al. 2014). Multiyear monitoring of sites, chosen as gateways or hotspots, would yield precise migration timing data and individual mortality assessments. When coupled with local resource availability sampling, this system could provide an unprecedented mechanistic view into how foraging animals utilize available resources.

The high temporal resolution of TOA tracking could additionally be used to answer questions about group dynamics, collective decision making and social information use. How are movements tied to weather, climate (Lyon et al. 2008), habitat manipulation and fragmentation (Sekercioglu 2007)? What conditions dictate the range of dispersal and where do these animals go? Though current tracking tools allow coarse migratory connectivity to be studied (Webster et al. 2002, Marra et al. 2010), better spatiotemporal data could reveal how migrants utilize specific resources at each stopover point. Are there critical stopover locations without which a migratory sub-population would be expected to collapse? How might reductions in the habitat quality at a staging area impact a migratory species (Piersma 2012)? Are there thresholds below which the resource is no longer viable? Answering questions like these will require large amounts of location data with high temporal and spatial resolution; TOA tracking systems, deployed at locations of interest, will make studies like these possible. Additionally, though we have highlighted applications to small, winged species, this technology is well suited to the study of larger animals, who are capable of carrying larger conventional tracking devices: inexpensive, long-lived tags enable large numbers of individuals to be tracked, low mass offers attachment flexibility, while the recurring costs to obtain data from the system are lower than competing techniques.

METHODS

The Real-time Location System that we have built employs mobile transmitters and a fixed network of time-synchronized receivers. The receiver network continuously ’listens’ for tag transmissions and when they are detected, the arrival time is precisely measured (+/30 nanoseconds). The arrival time, a unique tag identifier code, and additional metadata about the status of the receiver are sent from each detecting receiver to a central server where they are stored in a database as an ”event”. When the server identifies groups of events that are likely to have originated from the same tag transmission, an algorithm uses the arrival times to simultaneously estimate the tag position and transmission time. This system exploits two key engineering concepts: spread spectrum signals and matched filter detection.

The term spread spectrum refers to a class of methods of expanding the amount of radio-frequency bandwidth that is used to transmit a particular signal. For example, if a spread spectrum transmitter is tasked with broadcasting a message that occupies 10 Hz of bandwidth, it might use 100 Hz of bandwidth to actually send the message. Though this might seem wasteful, spread-spectrum approaches offer several advantages relative to narrow-band methods, including increased effective signal strength, improved interference rejection, and more accurate timing resolution. Our RTLS exploits these three properties, to respectively: increase tag reception range, allow multiple tags to operate simultaneously in the same region, and precisely measure the signal propagation time.

Each tag’s transmitter performs spread-spectrum modulation by multiplying (mixing) two signals together. The first is a fixed-frequency (typically150 to 450 MHz) sinusoid,

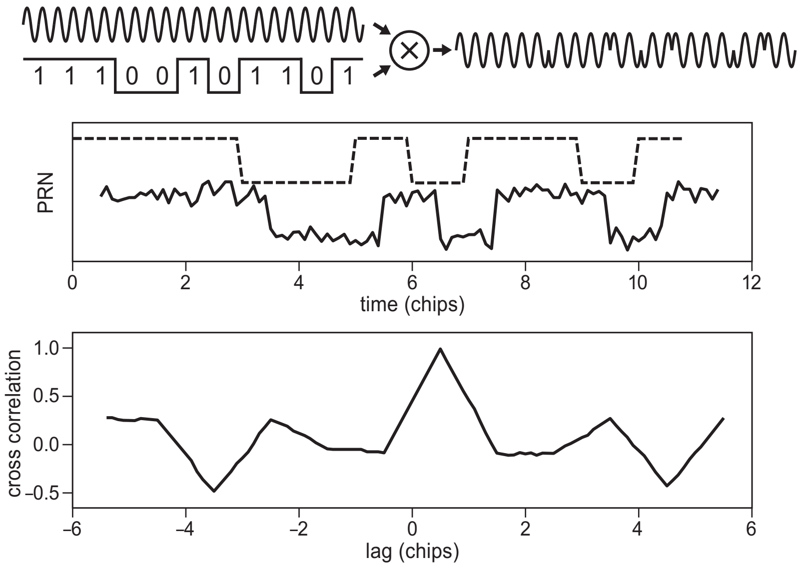

Figure 3.4 [Top]: Binary Phase-Shift Keying (BPSK) modulation of carrier by pseudo-random noise (PRN); [Middle]: PRN signal (dashed) with a noisy, delayed version; [Bottom]: Cross-correlation of the noisy delayed signal and the original PRN.

called the carrier; the second is a digital signal whose sequence of 1s or 0s is approximately random, but repeatable -a so called pseudo-random-noise (PRN) sequence; part of this sequence is common to all the tags and part of it is unique to each tag. In our tags, this multiplication has the effect of inverting the phase of the carrier in the modulated signal when a 1 is present, as illustrated at the top of Figure 3.4, a so-called Binary Phase-Shift Keying modulation scheme. This modulation is relatively easy to produce with discrete circuit building blocks (24).

$$R(l) = \sum_{m=-\infty}^{\infty} x^*(m)y(l+m) \tag{1}$$

This signal is demodulated at each in-range receiver by mixing it with another fixed-frequency sinusoid, restoring the PRN signal from the tag, albeit with additional noise. The middle chart in Figure 3.4 shows a simple example, with the dashed line representing the original PRN and the solid line representing the received PRN, corrupted by additive white Gaussian noise (AWGN) and a delay. Since the receiver knows all tag PRNs *a priori* it can search for the presence of the transmitted signal using a cross-correlation computation (equation 1). The cross-correlation R(l) of the two example signals is shown at the bottom of Figure 3.4. Notice that the domain of R(l) is the relative lag of the two signals and the range is the relative match between the two signals. The receivers use the peak cross-correlation value in combination with an adaptive threshold detection algorithm to determine that a PRN is received; the lag at that value provides an estimate of reception time.

Portability and flexibility were key design goals for this tracking system; permanently installed AC-power supplies, large fixed towers, and heavy equipment are not required. The receivers are sensitive enough to achieve 5 km reception range using 5 m telescoping pole towers with omni-directional VHF antennas. Each receiving station has a total power requirement of 25W, including secondary data radio links to communicate with the central server. This allows the receivers to be solar powered, with modest battery capacity for low-light and night-time operation.

Although each tag uses a unique orthogonal code, the current receiver design is capable of detecting only a single tag's transmission at any particular time. This causes other tag transmissions that overlap in time to be ignored (the first tag to transmit is detected). Although the tag transmissions are brief, relatively infrequent, and do not occur on exactly the same schedule (by design), if a large enough number of tags are within range of the same receiver, the receiver will inevitably fail to detect some tag transmissions. The percentage of missed tag transmissions as a function of the number of in-range transmitters and their transmission interval is shown in Figure 3.5. Note that this is not an intrinsic limitation of the method; this is a limitation of the current implementation. More powerful processors at the receivers will allow simultaneous tag detections.

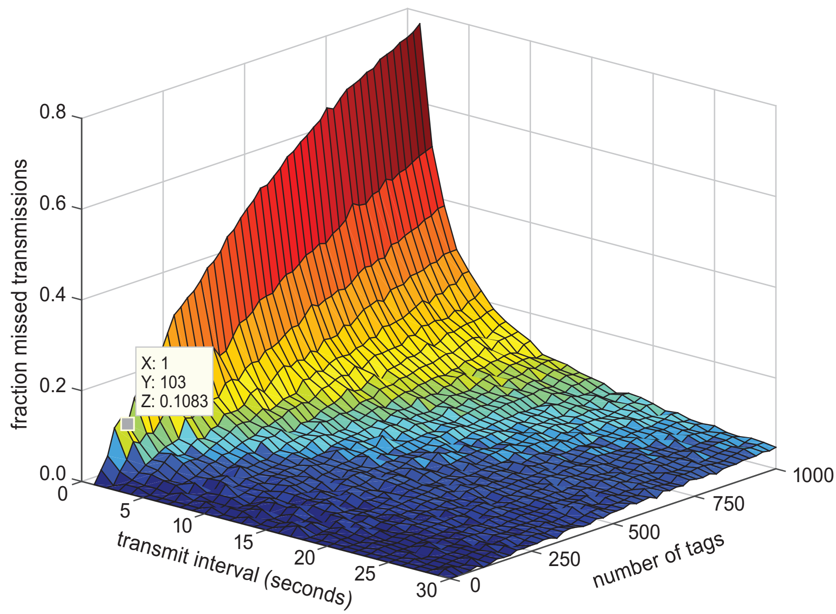

Figure 3.5 The current receiver design ignores simultaneous tag transmissions, therefore if enough transmitters are within range of a single receiver, some percentage of the total number of tag transmissions will be ignored. As the plot shows, a large number of tags can currently be accommodated even with transmit intervals shorter than half a minute. The cursor position shown corresponds to 103 in-range transmitters, each transmitting once per second. The current receivers will miss 11% of these transmissions. Future receivers, based on updated processors, will improve this capability.

When a receiver detects a tag transmission, it packages the tag id (determined by PRN number) along with the time of reception and other meta-data into a UDP datagram and sends the message via IP-radio equipment to a centralized server where the message is added to a database. Off-line operation is also possible; in this case the data from each receiver are stored locally and inserted into the database when field staff service the receiver. When the server identifies database entries from different receivers that are likely to have been from the same tag transmission event it attempts to compute a position estimate using a least-squares pseudo-range algorithm, similar to the method employed by GPS receivers. Position estimates from transmission locations within the receiver array have a 1σ error of 10 m. End-users may query the database remotely for tag position data. The frequency of position updates depends primarily on the frequency of tag transmissions, a user-specified parameter that can be traded-off against desired tag lifetime and mass to suit a particular application. Update rates as fast as 1 second are possible. Extensive technical details and performance measurements can be found in MacCurdy et al. (2008, 2009, 2012).

Section II

Sociality