Characterizing the Accuracy of a Self-Synchronized Reverse-GPS Wildlife Localization System

Adi Weller Weiser*, Yotam Orchan*, Ran Nathan*, Motti Charter*, Anthony J. Weiss†, and Sivan Toledo†
*The Hebrew University and †Tel-Aviv University, Israel

***Abstract*—We characterize the accuracy of a wildlife localization system that is based on the reverse-GPS or time-of-arrival (TOA) principle, in which radio receivers at known locations collaborate to determine the location of a transmitter attached to a wild animal. We describe the system and its real-world deployment in detail and show that it produces accurate location estimates in real settings and over long periods of time. Localization errors of wild animals carrying our transmitters (when stationary in nest boxes) have standard deviation of about 5m and mean of 5–15m (but sometimes as high as 24m). The mean error changes slowly with a time constant of hours or more and is frequency dependent; in the center of the receiver's band-pass filters, the mean error is usually about 5m. These characterizations are of raw localizations, without path-based filtering and smoothing. We also show how to reliably quantify the error in each individual location estimate. In addition to the characterization of accuracy, we introduce three technical innovations in the system. First, a method to model the error of individual time-of-arrival measurements, enabling correct weighing of the data to estimate locations and allowing estimation of the covariance matrix of each location estimate. Second, extensive use of known-position beacon transmitters, to synchronize the clocks of receivers (radio-frequency TOA localization require accurate clocks), to characterize and continuously monitor the performance of the system, and to model arrival-time-estimation errors. Third, we estimate the covariance matrix of each localization.**

I. Introduction

Although GPS technology has revolutionized wildlife tracking [2], [14], a strong need remains for wildlife tracking technologies capable of addressing missing or weak features in GPS tracking systems [22]. Several research projects attempted to develop such technologies [1], [3], [6], [9], [10], [11], [15], including the *reverse-GPS* technology [8], which estimates an animal's location from measurements of the time of arrival (TOA) of a radio signal emitted by a *tag* (transmitter) attached to the animal. This technology aims to track wildlife with higher throughput, lower cost, and lighter tag weights than is possible with GPS, but within one geographic region (tens of kilometers). These goals are similar to those of automated angle-of-arrival (AOA) wildlife tracking systems [6], [15], but AOA systems suffer from several disadvantages relative to TOA systems: accuracy degradation that is linear with distance from receivers, higher cost (multiple antennas and the circuitry to handle signals from them), and usually low throughput. At least in principle, reverse-GPS (TOA) better fits the needs of regional wildlife tracking.

This research was supported by the Minerva Center for Movement Ecology and by grants 965/15 and 863/15 from the Israel Science Foundation (funded by the Israel Academy of Sciences and Humanities).

Our main aim in this paper is to characterize and validate the accuracy of reverse-GPS wildlife tracking systems in real-life settings and with free-ranging animals. First, we characterize the typical accuracy that users can expect from such systems. Are location estimates usually accurate to within meters, tens of meters, or worse? The accuracy of individual location estimates varies with the geometry of transmitter and receivers and with signal-to-noise ratios (SNRs). Our second aim is to quantify the uncertainty in each location estimate using an estimated covariance matrix. A characterization of the uncertainty of each location estimates is useful both for constructing higher-level metrics of movements (e.g., smoothed paths, estimates of velocity and acceleration and detection of pauses) and for interpretation of location data by users (ecologists and biologists). This is particularly useful, though difficult, to characterize errors in location estimates that are "reasonable", as opposed to gross outliers that are relatively easy to detect. We do not consider detection of gross outliers in this paper because it can be done with general-purpose methods that are not specific to TOA techniques.

The paper also documents a new reverse-GPS wildlife tracking system called ATLAS, which we use as a platform for the investigation of accuracy, as well as its first large-scale long-term deployment. Figure 1 presents a sample output of the system.

ATLAS includes three significant advances over the pioneering reverse-GPS wildlife tracking system of MacCurdy et al. [8] and over automated AOA wildlife tracking systems [6], [15]. One is the extensive use of beacon transmitters at known locations. Beacon transmitters enable synchronization of the clocks of ATLAS's receivers, allowing the use of unsynchronized and fairly inaccurate clocks; this is a significant improvement over the reliance on absolutely accurate clocks in [8]. Both our system and [8] use GPS-disciplined clocks, but [8] depends on the absolute accuracy of the clock, whereas ATLAS depends mostly on the stability of the clock over time, which is easier to achieve. Beacons also allow us to characterize the overall accuracy of localizations and to monitor the performance of the system continuously.

A method to quantify the uncertainty in each arrival-time measurement from the SNR is the second novelty in ATLAS. This method uses a model that we fit to actual detection reports of beacons in the field. This quantification of uncertainty is crucial both to the correct weighting of detection reports, and to the construction of the covariance matrix of a location estimate. Our method overcomes a significant difficulty: the

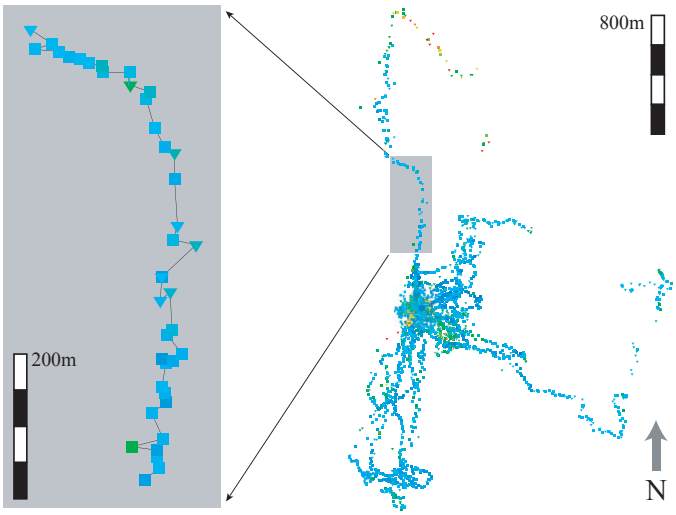

Figure 1. Localizations of one Barn Owl (number 243) during one week in May 2015. The section that the figure zooms in on represents 134s of flight. The squares and triangles represent raw localizations; squares denote localizations based on reception of the signal by 5 or more receivers, triangles by 4. The blue-to-green-to-red color reflects the norm of the covariance matrix of the localization (blue for small, then cyan and green, red for large). Quantitative error statistics of these localizations are presented in Figure 7 below.

direct measurement of the error of an individual arrival time estimate is essentially impossible. We address this by combining quadruplets of arrival-time estimates in a way that yields a linear combination of four errors with coefficients ± 1, from which we can estimate the variance of individual error.

Finally, ATLAS estimates locations using a maximum-likelihood method. Maximum-likelihood estimates are known to be asymptotically unbiased and optimal in the sense that the variance achieves the Cramer-Rao lower bound [21] (asymptotically refers in this context to a large number of measurements or high signal-to-noise ratios). We therefore prefer this estimation technique to the geometric techniques that have been used so far in AOA and TOA automated wildlife radio tracking systems. The maximum-likelihood approach also facilitates the construction of covariance matrices and/or confidence ellipses for each location estimate; none of the earlier systems [6], [8], [15] produced such outputs. We note that GPS receivers can estimate confidence ellipses of location estimates, which are essentially equivalent to covariance matrices, so wildlife tracking systems based on GPS can produce localization-quality metrics similar to ours. But these types of metrics have never been previously used in automated non-GPS wildlife tracking.

The rest of this paper is organized as follows. Section II explains the mathematical modeling of the TOA localization problem. Section III describes the overall design of ATLAS, focusing on the details that are pertinent to the accuracy of the system and to its evaluation. Section IV presents experimental results that characterize the accuracy of the system. We present our conclusions in Section V.

II. Mathematical Formulation and Algorithms

This section develops the location estimation models that ATLAS uses (Sections II-B and II-C) and the linear-algebra, calculus, and numerical-analysis tools that we use to solve the problems and to quantify the uncertainty in the solutions (Sections II-D–II-G). We begin by establishing notation in the following paragraphs and by explaining the intuition and engineering consideration behind the use of beacons to synchronize the clocks in the system (Section II-A). The presentation in Sections II-B and II-C is important for understanding some of the design decisions in ATLAS and some of our experimental results. Section II-B also presents an innovative and useful approximation technique for the case of inaccurate clocks. The techniques that we use in Sections II-C–II-G are fairly routine (see, e.g., [20]); we include them for completeness and reproducibility and because the results are not trivial to derive.

The localization algorithm processes N_r *detections* of N_t different transmissions from one unknown location and from multiple known locations by M receivers. A transmission is typically detected by more than one receiver, so $N_r > N_t$. We name (or index) the detections using the integers $1,\ldots,N_r$ and the transmissions using $1,\ldots,N_t$. Each detection $i \in \{1,\ldots,N_r\}$ is of some transmission $j \in \{1,\ldots,N_t\}$. We map a detection index i to the corresponding transmission index j using the notation $j(i)$. We order the detections so that the first ones are from the unknown location, and we denote the number of these detection by N_u. We denote the estimated arrival time of detection $i \in \{1,\ldots,N_r\}$ by t_i in the local clock of receiver i and by $\bar{t}_i$ in real time (relative to some fixed starting time). The arrival time is estimated from samples produced by the analog-to-digital converter (ADC) of the receiver. We denote the delay of the signal between the antenna and the ADC by g_i. We denote the (real time) unknown transmission time of $j(i)$ by $\tau_i = \tau^{(j(i))}$. We denote the index of the receiver of detection i by $r(i) \in \{1,\ldots,M\}$ and its known position by $p_i = \begin{bmatrix} x_i & y_i & z_i \end{bmatrix} = p^{(r(i))}$. We denote the index of the transmitter of transmission $j(i)$ by $b(i)$ and its position by $\pi_i = \pi^{(b(i))}$. Note that $\pi_1 = \pi_2 = \cdots = \pi_{N_u}$ and that this position is not known, whereas $\pi_{N_u+1},\ldots,\pi_{N_r}$ are known (there are also repetitions among them). Our notation uses Greek letters to denote transmission times and transmitter positions and Latin letters to denote arrival (reception) times and receiver positions.

The fundamental constraint associated with detection i is

$$\bar{t}_i = \tau_i + \tfrac{1}{c}\,\|\pi_i - p_i\|_2 + g_i + \varepsilon_i$$

where c is the speed of propagation (speed of light for radio signals) and ε_i represents the error in the measurement or estimation of the arrival time of the signal. We assume the errors have zero mean and that they are uncorrelated. If the receivers have synchronized perfect clocks so $t_i = \bar{t}_i$, then

$$t_i = \tau_i + \tfrac{1}{c}\,\|\pi_i - p_i\|_2 + g_i + \varepsilon_i\,.$$

A. Properties of Clocks that Drive System Design

The accuracy of a TOA localization system depends on its clocks. In reverse-GPS systems with digital (sampling) receivers, the relevant clocks are the sampling clocks that drive the ADCs. Each ADC sample can be associated with a particular tick of the sampling clock, and in some cases, the time of one sample is known to some degree of accuracy in an external time frame, for example coordinated universal time (UTC). If the clock rate does not vary and is known exactly, then we can relate all the samples to the external time frame.

One approach, which has been used by MacCurdy et al. [8], uses GPS-disciplined clocks [7] and assumes that they are perfectly accurate with respect to UTC. The assumption is not correct but it simplifies the localization system. A GPS-disciplined clock locks a crystal oscillator whose temperature is stabilized by an oven to the atomic clocks of the GPS system. These clocks achieve good clock-rate stability and accuracy; the ones that we use are accurate to within 1ppb [5], which is typical. That is, they err by about 1ns (equivalent to about 30cm at radio propagation speed) per second. These clocks can also associate a tick with UTC, but unfortunately this association is much less accurate: our clocks achieve $\pm$50ns RMS error, which is also typical. This implies that errors of 100ns or more (30m) occur about 5% of the time.

ATLAS uses a different approach, which is roughly equivalent to establishing an external time frame that starts at time $\bar{t}$ in a beacon transmitter at a known location π transmits a signal. A receiver at location p that receives the signal and estimates its arrival time to be t in the local clock can deduce that global time $\bar{t}$ coincides with local time $t - \|p-\pi\|/c$. This kind of reasoning can allow the system to correctly align detection reports from different receivers correctly.

We use this idea below but in a somewhat more sophisticate way that combines the synchronization of the clocks with the localization task and that takes into account the uncertainty in the clock synchronization.

B. Coping with Unsynchronized Clocks

If the clocks are not perfectly synchronized, the constraints become more complex. We assume that the clocks have arbitrary offsets and slightly incorrect rates. We also assume that the clock rate at each receiver remains essentially constant during the period that spans the receiver's detections. Let ℓ be some fixed detection at receiver $r(i)$. We assume that at real time $\hat{t}$, the clock of this receiver shows

$$\bar{t}_\ell + o^{(r(i))} + (1+k^{(r(i))})(\hat{t}-\bar{t}_\ell) = \hat{t} + o^{(r(i))} + k^{(r(i))}(\hat{t}-\bar{t}_\ell)$$

where $o_i = o^{(r(i))}$ is the offset of the clock of $r(i)$ at time $\bar{t}_\ell$ and $k_i = k^{(r(i))}$ is its clock-rate error. The constraint for detection i now becomes

$$t_i = \bar{t}_\ell + o_i + (1+k_i)\left(\tau_i + \tfrac{1}{c}\|\pi_i - p_i\|_2 + g_i - \bar{t}_\ell\right) + \varepsilon_i\,,$$

or if we use superscript to indicate the identity of most quantities,

$$\begin{aligned} t_i &= \tau^{(j(i))} + \tfrac{1}{c}\left\|\pi^{(b(i))} - p^{(r(i))}\right\|_2 + g_i + o^{(r(i))} \\ &\quad + k^{(r(i))}\left(\tau^{(j(i))} + \tfrac{1}{c}\left\|\pi^{(b(i))} - p^{(r(i))}\right\|_2 + g_i - \bar{t}_\ell\right) \\ &\quad + \varepsilon_i\,. \end{aligned}$$

In principle, if we assume that $g_i = g^{(r(i))}$ (the analog delay depends only on the identity of the receiver; we shall discuss this assumption below) we can use weighted least-squares minimization to obtain estimates of the unknown position π_1, the transmission times τ, the clock offsets o, and the clock rates k, and the analog delays g. However, this is a large non-linear least-squares problem that is likely to be difficult to solve. We can simplify it considerably by making one of two assumptions about the clock rates $1+k_i$. If we assume that they are exact so $k_i = 0$, then the only non-linearity that remains involves the distances between transmitters and receivers; we cannot hope to eliminate this non-linearity, but at least the dependence on the time unknowns τ, o and g is linear. Furthermore, under the assumption that $g_i = g^{(r(i))}$ we can estimate the sum $g+o$ rather than estimate both g and o. Alternatively, note that

$$\bar{t}_i = \tau_i + \tfrac{1}{c}\|\pi_i - p_i\|_2 + g_i + \varepsilon_i$$

or

$$\tau_i + \tfrac{1}{c}\|\pi_i - p_i\|_2 + g_i = \bar{t}_i - \varepsilon_i\,.$$

Since detections i and ℓ are at the same receiver, $o_i = o_\ell$ and $k_i = k_\ell$, so

$$t_i - t_\ell = (1+k_i)(\bar{t}_i - \bar{t}_\ell) + k_i(\varepsilon_\ell - \varepsilon_i) \tag{1}$$

. Therefore, the non-linear clock-rate-error term becomes

$$\begin{aligned} &k_i\left(\tau_i + \tfrac{1}{c}\|\pi_i - p_i\|_2 + g_i - \bar{t}_\ell\right) = \\ &\qquad k_i(t_i - t_\ell) - \frac{k_i^2}{1+k_i}(t_i - t_\ell) - \frac{k_i^2}{1+k_i}(\varepsilon_\ell - \varepsilon_i) - k_i\varepsilon_i\,. \end{aligned} \tag{2}$$

If we approximate the $k_i\left(\tau_i + \frac{1}{c}\|\pi_i - p_i\|_2 - \bar{t}_\ell\right)$ term in constraint i by $k_i(t_i - t_\ell)$, we replace a non-linear term by a linear one. The approximation increases the error in that constraint by at most

$$\begin{aligned} &\left| -\frac{k_i^2}{1+k_i}(t_i - t_\ell) - \frac{k_i^2}{1+k_i}(\varepsilon_\ell - \varepsilon_i) - k_i\varepsilon_i \right| \\ &\qquad \approx \left| -k_i^2(t_i - t_\ell) - k_i^2(\varepsilon_\ell - \varepsilon_i) - k_i\varepsilon_i \right| \\ &\qquad < k_i^2\,|(t_i - t_\ell)| + k_i^2\,|\varepsilon_\ell| + k_i^2\,|\varepsilon_i| + |k_i|\,|\varepsilon_i|\,. \end{aligned}$$

Free-running low-cost low-power temperature compensated crystal oscillators can achieve 10ppm ($|k_i| \le 10^{-5}$) accuracy over a wide range of temperature and several years of aging; oven-controlled oscillators can achieve slightly better accuracy and for longer periods, as well as better short-term stability. GPS-disciplined oscillators achieve accuracy of 1ppb or better [5], [7]. Therefore, we expect the magnitude of k_i to be at most 10^{-5}, which implies that the last term in our error

estimate above is likely to be at least 5 orders of magnitude smaller than ε_i, and the two middle terms are orders of magnitude smaller. The first term, $k_i^2\,|(t_i - t_\ell)|$, requires a more careful analysis.

With 10ppm clocks, the linearization adds about 10^{-10}s or 0.1ns of error for every second in the difference $t_i - t_\ell$. Therefore, if we want to limit the contribution of this error to about 1ns, we need to limit $t_i - t_\ell$ to about 10s. With a 1ppb GPS-disciplined clock, the linearization only adds 10^{-18}s of error for every second of difference $t_i - t_\ell$; even if we let $t_i - t_\ell$ span 1000s, the linearlization will contribute only a femtosecond (10^{-15}) to the error. Assuming that $k_i = 0$ can lead to much larger errors: 10μs if clock accuracy is only 10ppm and $t_i - t_\ell$ spans one second, or 1ns if clock accuracy is 1ppb and $t_i - t_\ell$ spans one second.

To summarize, our constraint system consists of constraints of the form

$$t_i = \tau^{(j(i))} + \tfrac{1}{c}\left\|p^{(r(i))} - \pi^{(b(i))}\right\|_2 + g_i + o^{(r(i))} + k^{(r(i))}\,(t_i - t_k) + \varepsilon_i\,. \tag{3}$$

We assume that $g_i = g^{(r(i))}$ and estimate the unknowns τ, $g+o$, k and π_1 by minimizing $\sum W_{ii}^2\varepsilon_i^2$ where W is a diagonal weighting matrix with $W_{ii} = \sigma^{-1}(\varepsilon_i) = \mathrm{var}^{-1/2}(\varepsilon_i)$, or an estimate of the inverse of the standard deviation. We refer to this estimate as a weighted least squares (WLS) estimate. Ignoring the approximation in the k_i term, WLS produces maximum-likelihood estimates if the error vector ε is distributed normally. Note that the only non-linear term in the constraint is the distance term. Clock offsets are relative to some reference time which is usually meaningless, so we can arbitrarily set one of the offsets to zero, say at the first receiver, $o^{(1)} = 0$. Clock rates are similarly relative so we can set $k^{(0)} = 0$.

If we assume that $k = 0$, 5 detections of a transmission from an unknown location and 5 detections of a beacon transmission (by the same 5 receivers) are sufficient to estimate the 3-dimensional unknown location π, the 4 offset+delay elements of $g+o$ (one is set to zero), and the two transmission times. If the z value is known (or otherwise fixed), 4 detections from the unknown locations and 4 from a beacon are sufficient. To estimate k, we need detections of one more beacon transmission. Detections from additional receivers and detections of additional beacon transmission provide redundancy that helps reduce estimation errors.

C. A Simplified Version

A simplified version can be used for generating sub-optimal localizations quickly. These can be used for real-time localization if the computational load of the full WLS problem is too high, or for quickly generating good initial guesses for the non-linear solver that handles the full problem.

In this version we assume that the set of detections includes only one transmission from an unknown location and one transmission from a known location, and we assume that $k_i = 0$. This leads to a system of pairs of constraints of the form

$$\begin{aligned} t_i &= \tau^{(1)} + \tfrac{1}{c}\left\|\pi^{(1)} - p^{(r(i))}\right\|_2 + g_i + o^{(r(i))} + \varepsilon_i \\ t_{i'} &= \tau^{(2)} + \tfrac{1}{c}\left\|\pi^{(2)} - p^{(r(i))}\right\|_2 + g_{i'} + o^{(r(i))} + \varepsilon_{i'} \end{aligned}$$

We paired constraint i at receiver $r(i)$ involving the unknown location $\pi^{(1)}$ with constraint i' at the same receiver $r(i) = r(i')$ involving the known location $\pi^{(2)}$. Instead of minimizing the sum of squares, substitute

$$o^{(r(i))} = t_{i'} - \tau^{(2)} - \tfrac{1}{c}\left\|\pi^{(2)} - p^{(r(i))}\right\|_2 - g_{i'} - \varepsilon_{i'}$$

in the ith constraint to obtain

$$\begin{aligned} t_i &= \tau^{(1)} - \tau^{(2)} + \tfrac{1}{c}\left\|\pi^{(1)} - p^{(r(i))}\right\|_2 + g_i \\ &\quad + t_{i'} - \tfrac{1}{c}\left\|\pi^{(2)} - p^{(r(i))}\right\|_2 - g_{i'} + \varepsilon_i - \varepsilon_{i'} \end{aligned}$$

or assuming that $g_i = g_{i'}$ (again these are two delays at the same receiver),

$$\begin{aligned} t_i - \left(t_{i'} - \tfrac{1}{c}\left\|\pi^{(2)} - p^{(r(i))}\right\|_2\right) \\ = \tau^{(1)} - \tau^{(2)} + \tfrac{1}{c}\left\|\pi^{(1)} - p^{(r(i))}\right\|_2 + \varepsilon_i - \varepsilon_{i'}\,. \end{aligned}$$

We now minimize the (weighted) sum of squares of $(\varepsilon_i - \varepsilon_{i'})$ to estimate the unknown $\pi^{(1)}$ and the difference $\tau^{(1)} - \tau^{(2)}$. As we shall see below, this minimization problem is considerably easier than the full problem. It requires detection of a transmission from an unknown location and from a beacon by 4 receivers to estimate a 3-dimensional $\pi^{(1)}$ and $\tau^{(1)} - \tau^{(2)}$, or 3 receivers if the z coordinate of $\pi^{(1)}$ is known.

D. Useful Transformations of the Minimization Problems

Certain manipulations of the non-linear least squares problems that the system solves to estimate p can make them easier.

Eliminating Linear Terms: All the problems that we have discussed are of the form $\min_{\xi,\zeta}\|\beta - \phi(\xi,\zeta)\|_2$ where β is a known vector, ξ and ζ are vectors of unknowns (sometimes one-dimensional) and ϕ is a non-linear vector function of ξ and ζ. Furthermore, ϕ is always separable: it splits into a nonlinear and a linear term, $\phi(\xi,\zeta) = \nu(\xi) + L\zeta$, where ν is a non-linear vector function and L is a matrix.

We can eliminate the unknowns ζ that the residual depends on linearly. Let $L = QR$ which Q being an orthonormal basis for range(L), which implies that R is square and invertible. We can obtain Q and R from a QR factorization or a singular value decomposition (SVD) of L; in both cases R^{-1} is easy to apply. For a given ξ,

$$\min_\zeta \|\beta - \phi(\xi,\zeta)\|_2 = \min_\zeta \|\beta - \nu(\xi) - L\zeta\|_2 \tag{4}$$

is a linear least squares problem with a minimizer $\zeta^\star = R^{-1}Q^T(\beta - \nu(\xi))$ and a minimum

$$\|\beta - \nu(\xi) + L\zeta^\star\| = \left\|\left(I - QQ^T\right)(\beta - \nu(\xi))\right\|_2\,. \tag{5}$$

Therefore,

$$\min_{\xi,\zeta} \|\beta - \phi(\xi,\zeta)\|_2 = \min_{\xi} \left\| \left(I - QQ^T\right)(\beta - \nu(\xi)) \right\|_2 .$$

The non-linear minimization problem on the right has a lower dimension (only the unknowns in ξ) is therefore likely be easier to solve than the problem on the left. Once the minimizer $\xi^\star$ is found, we can easily compute $\zeta^\star = R^{-1}Q^T(\beta - \nu(\xi))$ if it is needed.

In our problems, $\xi = p$, the sought-after position variable, and ζ consists of transmission times or differences of transmission times, clock offsets, and possibly clock-rate errors; in general we do not need $\zeta^\star$, but it is clearly possible and inexpensive do to so.

Note that the orthogonal projector $I - QQ^T$ can also be expressed as

$$I - L\left(L^TL\right)^{-1}L^T = I - QQ^T . \tag{6}$$

The computational cost of $I - QQ^T$ and that of applying $I - L\left(L^TL\right)^{-1}L^T$ are similar, but the former expression is more numerically stable. Also, in many cases we can form Q and $I - QQ^T$ symbolically, which further simplifies the computations. We discuss these important special cases below.

In the weighted case, W is a diagonal matrix with diagonal entries and we minimize

$$\min_{\xi,\zeta} \|W(\beta - \phi(\xi,\zeta))\|_2 = \min_{\xi} \| \left(I - \tilde{Q}\tilde{Q}^T\right) W(\beta - \nu(\xi)) \|_2 \tag{7}$$

where $\tilde{Q}$ is an orthonormal basis for range(WL).

1) Scaling and Shifting: Some numerical algorithms, including some linear and non-linear least squares solvers, are sensitive to large dynamic range in their inputs and outputs. The dynamic range is sometimes an artificial artifact of the choice of units and origin. To reduce the risk of large dynamic range and the resulting numerical instability, we scale quantities so that distances and locations are specified in meters and time differences in nanoseconds. This yields propagation speed $c \approx 0.3$m/ns and puts distance and time-difference quantities on similar scales. In contrast, using SI units of meters and seconds leads to $c \approx 3 \times 10^8$, making time-difference quantities several orders of magnitude smaller than distance quantities. We also shift the time and position coordinates so that the results of localization problems are close to the origin, to avoid values that are unnecessarily large.

E. The Jacobian, Gradient, and Optimality Conditions

Many optimization algorithms require the gradient of the function that we minimize. The gradient is also helpful for testing for optimality once an approximate minimum is found, even if the minimization algorithm only uses the value of the function and not its gradient.

Our position estimate is

$$\arg\min_{\pi} \|A(\beta - \nu(\pi))\|_2^2 = \arg\min_{\pi} \|f\|_2^2 = \arg\min_{\pi} f^T f$$

where β is a known real vector, $\pi = \begin{bmatrix} x & y \end{bmatrix}^T$ or $\pi = \begin{bmatrix} x & y & z \end{bmatrix}^T$ is an unknown position in two or three dimensions, ν is a non-linear vector function of π, and A is a matrix that represents a projection and weighting. The structure of ν is very simple. If constraint i corresponds to a beacon detection, then $\nu_i = 0$. Otherwise,

$$\nu_i = \frac{1}{c}\|p^{(r(i))} - \pi\|_2 = \frac{1}{c}\sqrt{\left(x - x^{(r(i))}\right)^2 + \left(y - y^{(r(i))}\right)^2 + \left(z - z^{(r(i))}\right)^2} .$$

In this case,

$$\frac{\partial \nu_i}{\partial x} = \frac{2x - 2x^{(r(i))}}{2c\|\pi - p^{(r(i))}\|_2} = \frac{x - x^{(r(i))}}{c\|\pi - p^{(r(i))}\|_2}$$

and similarly for y and z. Computing ν and $\partial\nu/\partial\pi$ costs $O(N_u)$ arithmetic operations.

The gradient of the objective function is

$$\frac{\partial}{\partial \pi}\left(f^T f\right) = -2\left(\beta - \nu(\pi)\right)^T A^T A \frac{\partial}{\partial \pi}\nu(\pi) . \tag{8}$$

Note that since $A = \left(I - \tilde{Q}\tilde{Q}^T\right)W$ with both factors symmetric and $\left(I - \tilde{Q}\tilde{Q}^T\right)$ being a projection,

$$A^T A = WA . \tag{9}$$

We can also derive the gradient using the Jacobian,

$$\frac{\partial}{\partial \pi}\left(f^T f\right) = 2f^T \frac{\partial f}{\partial \pi}$$

At a minimum of the objective function, its gradient is zero. We use this both to test for optimality and to determine the sensitivity of the position estimate to timing errors.

The Jacobian of f is

$$\frac{\partial f}{\partial \pi} = -A\frac{\partial \nu}{\partial \pi} . \tag{10}$$

If we plug this into the computation of the gradient, we obtain

$$\begin{aligned} \frac{\partial}{\partial \pi}\left(f^T f\right) &= 2f^T \frac{\partial f}{\partial \pi} \\ &= -2\left(\beta - \nu(\pi)\right)^T A^T A \frac{\partial \nu}{\partial \pi} . \end{aligned}$$

F. The Covariance Matrix of a Localization

At a point that zeros the gradient of $f^T f$, the covariance matrix of the estimate π is

$$\left(\frac{d\pi}{dt}\right) W^{-2} \left(\frac{d\pi}{dt}\right)^T \tag{11}$$

where W^{-2} is the covariance matrix of the arrival times vector t and $d\pi/dt$ is the total-derivative matrix that relates infinitesimal timing displacements to infinitesimal position displacements (we use d to denote total derivatives and ∂ to denote partial derivatives). That is, if a measured arrival time t_i shifts to $t_i + \Delta t_i$, then the optimal solution π moves from from $\begin{bmatrix} x & y & z \end{bmatrix}^T$ to $\begin{bmatrix} x+\Delta x & y+\Delta y & z+\Delta z \end{bmatrix}^T$. We

apply the implicit function theorem to find the expression for $d\pi/dt = \begin{bmatrix}\Delta x & \Delta y & \Delta z\end{bmatrix}^T$. We have

$$\frac{\partial}{\partial \pi}\left(f^T f\right) = \begin{bmatrix}0 & 0 & 0\end{bmatrix} .$$

We transpose and take the total derivative of both sides with respect to t,

$$\frac{d}{dt}\left(\frac{\partial}{\partial \pi}\left(f^T f\right)\right)^T = 0_{3\text{-by-}m} . \tag{12}$$

By expanding the left-hand side using the chain rule and isolating $d\pi/dt$ we obtain

$$\frac{d\pi}{dt} = -\left(\frac{\partial}{\partial \pi}\left(\frac{\partial}{\partial \pi}\left(f^T f\right)\right)^T\right)^{-1} \frac{\partial}{\partial t}\left(\frac{\partial}{\partial \pi}\left(f^T f\right)\right)^T . \tag{13}$$

The 3-by-m left hand side is a product of the inverse of a 3-by-3 matrix and a 3-by-m matrix.

We now develop expressions for these matrices. We limit ourselves to the case in which A is independent of t (in formulations that estimate the clock rate this is not strictly true, although it is probably a reasonable approximation). We note that $\beta = \beta(t)$ and that ν is always independent of t. We have

$$\frac{\partial}{\partial t}\left(\frac{\partial}{\partial \pi}\left(f^T f\right)\right)^T = -2\left(\frac{\partial \nu}{\partial \pi}\right)^T A^T A \frac{\partial \beta(t)}{\partial t} . \tag{14}$$

Note that $\partial\beta/\partial t = I_{m\times m}$ because either $\beta_i = t_i$ or $\beta_i = t_i - \|\pi^{(b(i))} - p^{(r(i))}\|/c$ where $\|\pi^{(b(i))} - p^{(r(i))}\|/c$ is the distance between beacon $b(i)$ and receiver $r(i)$[1]. Therefore,

$$\frac{\partial}{\partial t}\left(\frac{\partial}{\partial \pi}\left(f^T f\right)\right)^T = -2\left(\frac{\partial \nu}{\partial \pi}\right)^T A^T A .$$

We now turn our attention to the other matrix in the expression for dp/dt. We note that ν is a function of p but β is not. If we derive by the vector π we obtain a product of a tensor and a vector; to avoid tensor notation, we compute this matrix by columns. The first column of

$$\frac{\partial}{\partial \pi}\left(\frac{\partial}{\partial \pi}\left(f^T f\right)\right)^T \tag{15}$$

is

$$\frac{\partial}{\partial x}\left(\frac{\partial}{\partial \pi}\left(f^T f\right)\right)^T = -2\frac{\partial}{\partial x}\left(\frac{\partial \nu}{\partial \pi}\right)^T A^T A(\beta - \nu) + 2\left(\frac{\partial \nu}{\partial \pi}\right)^T A^T A \frac{\partial \nu}{\partial x} . \tag{16}$$

The first term is a product of the 3-by-m matrix $\frac{\partial}{\partial x}\left(\frac{\partial \nu}{\partial \pi}\right)^T$ with the m-vector $A^T A(\beta - \nu)$, which gives a 3-by-1 column vector. The second term is the product of the 3-by-m matrix $(\partial\nu/\partial\pi)^T$, the m-by-m matrix $A^T A$, and the m-vector $\partial\nu/\partial x$.

[1]In the simplified version β is slightly more complicated; we omit the details due to lack of space.

The y and z columns of 15 are obtained in a similar way. The first row of $\frac{\partial}{\partial x}\left(\frac{\partial \nu}{\partial \pi}\right)^T$ is

$$\frac{\partial^2 \nu_i}{\partial x^2} = \frac{1}{c}\frac{\delta_i^2 - \left(x - x^{(r(i))}\right)^2}{\delta_i^3} \tag{17}$$

for is that correspond to unknown locations and zero for is that correspond to beacon transmissions. The second row (and similarly, the third) is

$$\frac{\partial^2 \nu_i}{\partial y \partial x} = \frac{1}{c}\frac{-\left(x - x^{(r(i))}\right)\left(y - y^{(r(i))}\right)}{\delta_i^3} \tag{18}$$

for i that correspond to unknown locations and zero beyond.

G. Computational Complexity

To find a least-squares minimizer $\pi^\star$, the minimization algorithm computes the residual and perhaps the Jacobian or gradient several times. Once the minimizer is found, we normally also compute the covariance matrix and sometimes the linear model ζ. We now estimate the computation complexity of computing the residual of the least-squares problem at a given point π, as well as the cost of computing the Jacobian, gradient, and covariance matrix. Using the notation of (4) we have

$$\beta_i = \begin{cases} t_i - \frac{1}{c}\|\pi^{(b(i))} - p^{(r(i))}\| & \text{if } i \le N_u \\ t_t & \text{otherwise,} \end{cases}$$

so the cost of computing β for a given π is $\Theta(N_u) = O(N_r)$.

The structure of L and the cost of computing an orthonormal basis U for its column space are as follows. The matrix L has N_r rows and $2(M-1)+N_t$ columns if we estimate clock rates and $M-1+N_t$ columns otherwise. To enable efficient computation of U, we order the columns of L starting with columns that correspond to offset unknowns $o_i = o^{(r(i))}$, followed by rate unknowns $k_i = k^{(r(i))}$, and finally the transmission time unknowns $\tau_i = \tau^{(j(i))}$. The matrix L is sparse. The sparsity structure of WL is similar. The first $M-1$ columns (corresponding to clock offsets) are already orthogonal to each other; the corresponding columns of U are scaled copies of these columns. The next $M-1$ (clock-rate offsets) are also orthogonal to each other, but not to all the first $M-1$. But each clock-rate column is orthogonal to the offset columns of the other receivers, so to orthogonalize the first $2(M-1)$ columns, all we have to do is to subtract the projection of column j from column $M-1+j$. Subtracting each projection costs $O(N_t)$ arithmetic operations because both column j and column $M-1+j$ contain only N_t non-zeros. Note that the cost bound is valid whether we use a Gram-Schmidt-type projection or a Householder reflections. In general, the transmit-time columns (the last N_t columns) fill when we orthogonalize L. Subtracting the projection of the first $2(M-1)$ columns cost $O(MN_t^2)$. The subsequent orthogonalization of the last N_t columns costs $O(N_r N_t^2)$ arithmetic operations. The dominant cost is the last one. The following theorem summarizes the discussion so far.

Theorem 1. *Forming the sparse matrix WL costs $O(MN_t+N_r)$ operations. Computing its QR factorization costs $O(MN_t+N_rN_t^2)$. The Q factor is sparse with $2(M-1)$ columns with only N_t nonzeros per column and with N_t dense columns. Applying the Q factor or its transpose to a vector costs $O(MN_t+N_rN_t)$ operations. Computing the SVD of WL costs $O(MN_t+N_rN_t^2+(M+N_t)^3)$.*

If we limit N_t to a constant (the minimum is 3 if we estimate clock rates and 2 if we do not), the asymptotic cost of computing L and its QR factorization and the cost of applying Q or its transpose drops to $O(M+N_r)=O(N_r)$, which is linear in both the number of detections and the number of receivers. The cost of computing the SVD of L is not linear even in this simple case because of the M^3 factor.

The cost of applying A and A^TA is about twice and four times the cost of applying the Q factor of WL. This implies that the cost of computing the residual of the least-squares problem, the sum of squares, the gradient, Jacobian and covariance matrix are all asymptotically the same, $O(MN_t+N_rN_t)$ in the general case and $O(N_r)$ if N_t is constant.Applying both W and $I-ww^T/w^Tw$ to a vector requires only $O(n)$ arithmetic operations, so applying A^TA also costs the same asymptotically.

We note that computations involving the simplified version are even cheaper, but we omit the details due to lack of space.

III. ATLAS System Design

Transmitters: ATLAS uses miniature frequency-shift-keying (FSK) transmitters (tags) that transmit 8192-bit codes at 1Mb/s, ±380kHz deviation, 10mW, around 434MHz [19] (Toledo et al. [19] describe our first-generation tags; our current tags use 0201-size passives and 4-layer boards and are slightly smaller). The tags use a CC1101 integrated crystal-controlled FSK transceiver and an MSP430 microcontroller. Complete tags with coating, batteries and antennas weigh 1g and up, depending on the batteries used. The tags consume about 30mA when transmitting. We powered tags with a single lithium coin cell, a lithium-thionyl chloride cell, a pair of zinc-air cells, or a pair of silver-oxide cells [17]. Tags with zinc-air size 10 cells that transmit every second last about 10 days and weigh about 1.5g; similarly-configured tags with lithium CR2477 or lithium-thionyl TL4902 cell last about 100 days and weigh 10g. Tags that transmit every 8s and are powered by silver oxide size 337 cells weigh about 1g and last about 5 days. Each tag transmits 1 or 2 unique pseudo-random codes repeatedly, one on each center frequency (channel). The repetition rate is accurate (to within 30us or so, and usually much less) and known to receivers, so once a receiver detects a transmission from a particular tag, it knows when to expect the next transmissions and it performs the required signal processing only on windows of samples that are known to contain a transmission. Repetition rates range from 2Hz to 0.125Hz.

Receivers: The receivers consist of a USRP N200 sampling radio with a WBX RF daughter board, a custom front-end unit [18], and a commercial mast-mounted low-noise amplifier (LNA) that is mounted close to the antenna. The LNA establishes a low system noise figure. The front-end unit filters out-of-band interference, amplifies the signal to maintain noise figure, and limits power to protect the receiver. For the discussion here it suffices to say that the system noise figure is about 1dB and that the bandwidth of the receiver is about 5MHz with very sharp skirts (this entails a tradeoff that we will discuss later). We use three different commercial antenna types, depending on what we felt was most appropriate for the location: vertically-polarized omni-directional collinear antennas, circularly-polarized Yagi antennas, and circularly polarized egg-beater antennas.

Limiting Analog-Delay Variations: The analog delay in a receiving station depends on many factors including cable lengths and types and on the delay in filters, amplifiers, and other analog signal processors. In some of these elements, such as cables, the delay is almost independent of frequency and temperature. In other elements, especially filters, the delay is highly dependent on frequency, and to the extent that temperature affects frequency response, it also affects the delay. We have attempted to control delay variations in two main ways. First, we include in each localization problem only transmissions at a single nominal frequency (or channel) with a single symbol rate and deviation. This aims to equalize the delay in all the detection reports. Second, we limit the time span of transmissions in a localization problem to a short period (normally a few seconds to a couple of minutes) in order to limit the variation in delay due to temperature changes in the receiver's filters. These actions aim to make g_i constant at every receiver, so we can replace g_i by $g^{(r(i))}$ which allows us to combine g with the offset o.

However, as we shall see in the experimental results below, it appears that g_i is not entirely constant at each receiver, even within a localization problem. We believe that the variation is mostly due to temperature-induced frequency drift at the transmitters, which causes the frequency of the tag that we localize and the frequency of the beacon(s) that we use to synchronize the clocks to differ slightly. These frequency differences can result in different delays for the two (or more) signals, because we use filters with a significant delay and a significant delay variation, ranging from 150ns in the middle of the passband to around 250ns at the edges [18].

Signal Processing: Samples from a receiver are collected and processed by a computer running Linux. We sample at 8.33MS/s. Processing begins with a soft FSK detector that generates a real sampled signal s with values in $[-1,1]$. We compute the correlation of s with a pattern signal p that we compute by generating samples of the transmitted signal and passing them through the same soft FSK detector. We correlate real rather than complex signals because our transmitters suffer an arbitrary phase shift when switching between the two FSK frequencies. We estimate the arrival time of the signal by finding the maximum-magnitude sample in the correlation signal and interpolating a quadratic that passes through it and through the two adjacent samples. The extremum of this parabola is the estimate of the arrival time. The estimate is a

sub-sample value (8.33MS/s samples are spaced 120ns apart). We also compute an estimate of the signal-to-noise (SNR) ratio of the received signal. The estimate that we use is

$$\frac{\|\text{signal}\|_2}{\|\text{noise}\|_2} = \frac{p(n)^T s(n+\delta)}{\sqrt{\|s(n+\delta)\|_2^2 - (p(n)^T s(n+\delta))^2}} \tag{19}$$

where $p(n)$ is the finite-length m-dimensional sampled pattern signal and $s(n+\delta)$ is the m consecutive samples of the received signal that correlate best with $p(n)$. The numerator represents the energy of s in the subspace spanned by p, which is a sum of the energy received from the transmitter and noise and interference energy in the same subspace. The denominator represents the total received energy minus the energy in the pattern subspace. At high and moderate SNRs, this expression is an excellent approximation of the true SNR, because almost all the energy in the pattern subspace comes from the transmitter. At very low SNRs, this approximation overestimates the true SNR. With a random Gaussian noise, this becomes significant at SNRs around $1/m$ (that is, *very* low).

Localization Algorithms: Reports of detections of tags are sent to a server using the Internet (remote receivers have cellular modems). The server groups detections of the same transmission, bundles detections of a transmission of a tag at an unknown location with detections of beacon transmitters at about the same time, and sends them to a localization engine. The localization engine employs four different minimization algorithms. The localization process starts with building one or two objective functions. One always corresponds to the simplified formulation in Section II-C; it is used to compute a rough location estimate. The other, which is used for the output location estimate, is either the same (this is used in real time for speed), or a WLS formulation using detections of either one beacon transmission (for speed) or of all the beacons within a time window. The WLS formulations, and especially the multiple-beacons one, are used to produce accurate localizations off line. Both the simplified and the WLS objective functions can have 2 or 3 spatial unknowns, depending on how we configure the software. Even in the 3D case, the algorithm reverts to 2D if not enough receiver received a transmission. Two-dimensional localizations use a fixed z value that is configured manually according to the local topography.

The non-linear minimization algorithms that ATLAS employs are a gradient-free Nedler-Mead search [13] that is used when the last estimated location of the tag has a high residual in the objective function, a simple exhaustive grid search with a 100m grid that is used when even the Nedler-Mead algorithm fails to find a reasonable solution, and a Levenberg–Marquardt gradient-based algorithm [12] that is used to converge to a local minimum. The three algorithms are first applied to the simplified objective function. If the run aims to find a WLS solution, it uses the simplified estimate as an initial guess for a single run of Levenberg–Marquardt on the WLS objective function.

Estimating the Variance of Arrival-Time Estimates: To weigh the localization constraints and to estimate the covariance of the optimal solution, we need to estimate W, the covariance matrix of the arrival-time estimates. We assume that it is diagonal and estimate the diagonal entries, which are variances of arrival-time estimates. We estimate the variance using the expression

$$\text{var}(t_i) = \alpha \frac{\|\text{noise}\|_2^2}{\|\text{signal}\|_2^2} + \beta\ , \tag{20}$$

where α and β are parameters that we estimate and $\|\text{noise}\|_2/\|\text{signal}\|_2$ is the inverse of the SNR estimate (19). This expression is based on the assumption that the errors are due to two sources: quantization errors with variance β, and intrinsic errors associated with extracting a parameter (arrival time) from a signal passing through a noisy channel. The intrinsic errors are bounded from below by the Cramer-Rao bound, which in this case has the form $\alpha_{\text{Cramer-Rao}}\|\text{noise}\|_2^2/\|\text{signal}\|_2^2$. The expression that we use for the variance of arrival times is based on the assumption that the errors behave like the Cramer-Rao bound, perhaps with a larger coefficient.

We estimate α and β from differences of time-differences of arrivals, or *TDOA differences* for short. Let τ_1 and τ_2 two transmissions times of the same beacon at position π, and let $t_{1,1}, t_{1,2}$ be arrival times estimates of the signal transmitted at τ_1 at receivers 1 and 2 with clock offsets o_1 and o_2 at positions p_1 and p_2, and similarly for $t_{2,1}$, $t_{2,2}$. We ignore clock-rate errors. We have

$$\begin{aligned}
t_{1,1} &= \tau_1 + \frac{1}{c}\|\pi - p_1\|_2 + o_1 + \varepsilon_{1,1} \\
t_{1,2} &= \tau_1 + \frac{1}{c}\|\pi - p_2\|_2 + o_2 + \varepsilon_{1,2} \\
t_{2,1} &= \tau_2 + \frac{1}{c}\|\pi - p_1\|_2 + o_1 + \varepsilon_{2,1} \\
t_{2,2} &= \tau_2 + \frac{1}{c}\|\pi - p_2\|_2 + o_2 + \varepsilon_{2,2}\ .
\end{aligned}$$

The TDOA difference of these arrival times cancels out all the quantities except for the errors,

$$(t_{2,2} - t_{1,2}) - (t_{2,1} - t_{1,1}) = \varepsilon_{2,2} - \varepsilon_{1,2} - \varepsilon_{2,1} + \varepsilon_{1,1}$$

and if we assume that the errors are uncorrelated,

$$\text{var}\,(t_{2,2} - t_{1,2} - t_{2,1} + t_{1,1}) = \sum_{i,j\in\{1,2\}} \text{var}(\varepsilon_{i,j}) = \alpha\left(\frac{\|\text{noise}_{2,2}\|_2^2}{\|\text{signal}_{2,2}\|_2^2} + \cdots + \frac{\|\text{noise}_{1,1}\|_2^2}{\|\text{signal}_{1,1}\|_2^2}\right) + 4\beta\ . \tag{21}$$

To estimate α and β, we compute the error and the sum of noise-to-signal ratios (NSRs) of a large number of TDOA differences. We choose detections with τ_1 close to τ_2 to minimize the effect of the clock-rate differences, which we do not model in this procedure. We group the errors by NSR bins, compute the sample variance in each bin, and estimate α and β using a non-negative least-squares fitting. We weight

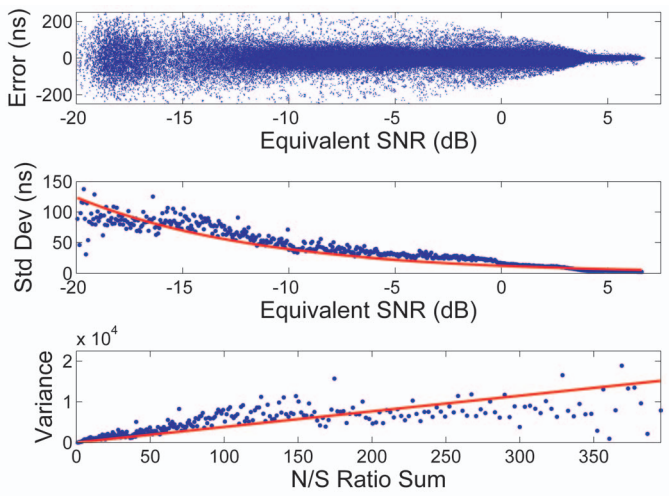

Figure 2. Estimating the variance of arrival-time errors. The individual TDOA differences are shown in the top scatter plot. The sample standard deviations and variances of errors in 0.05dB bins are shown in the bottom two graphs. The red curves show the weighted non-negative least-squares fitting of a line through the sample variances.

the least-squares problem by the number of samples in each bin. Figure 2 shows the results of this procedure when carried out on 4 hours of detections of one beacon by the 7 receivers deployed system described in the next section.

IV. Experimental Evaluation

We have deployed an ATLAS system with 7 receiving stations and 3 beacons in the flat, rural, and wildlife-rich Hula Valley in northern Israel [4], as shown in Figure 3. The absolute positions of receiving antennas and of beacon antennas were determined directly or indirectly by licensed surveyors. Five of the receiving antennas are placed on towers owned by a cellular company; the locations of these antennas were derived from the construction permits of the towers, with accuracy of about $\pm$1m. The location of two other receiving antennas and of beacon antennas were determined to within less than 1m by a surveyor that we hired.

This section mainly aims to characterize the accuracy of the system. We analyze 2-dimensional localizations based on the simplified formulation of Section II-C (except in Figure 8 below). The system automatically filters out localizations with a large residual $\|\varepsilon\|_\infty > 100$m and localizations that are more than 25km from the center of the region. In the analyses below, we also filtered localizations based on only 3 reception reports, which can be inaccurate and whose residual, which is always zero, provides no indication of the error. Fewer than 1% of the localizations of beacons were based on 3 reception reports. Of the localizations of two Barn Owls that we tracked, fewer than 10% were based on 3 reception reports, and of the localizations of a third owl, less than 22%.

We begin with a short overview of animal movement data that we have collected.

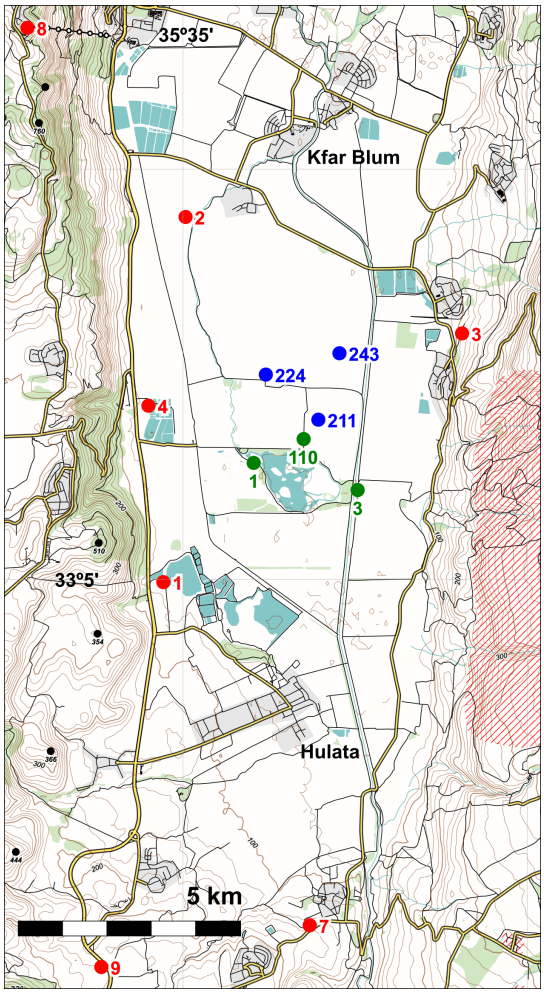

Figure 3. ATLAS receivers (red dots) and beacons (green dots) in the Hula Valley in northern Israel. The geodetic coordinates of beacon 1 are 33.106987N 35.600446E. The map also shows the location of three nest boxes of Barn Owls that we tracked (blue dots). The coverage of the system extends beyond this map; we collected useful localizations from an area of at least 15 by 40km. Map data © OpenStreetMaps, SRTM.

A. Scope of Overall Data Collection

Between April 2014, when we began tagging wild animals, and August 2015, we tagged animals of 9 species in the Hula Valley, mostly birds and bats. In these preliminary trials, some tags failed completely, some were received by too few receivers to be localized, and some failed prematurely. Problems included hardware malfunctions, firmware bugs, and attachment failures. Still, we collected more than 39 million localizations of 131 tagged animals of all 9 species, including 51 individuals with $> 100,000$ localizations and 6 with $> 1,000,000$. The largest number of individuals tracked simultaneously during a 24-hour period is 34 (greater mouse-tailed bats carrying 1.5–2.0 tags) on August 12, 2015, out of 40 bats that we tagged that day.

B. Statistics of Localization Errors

Figure 4 shows typical localizations of a fixed-position transmitter (beacon). Localizations at each channel form a

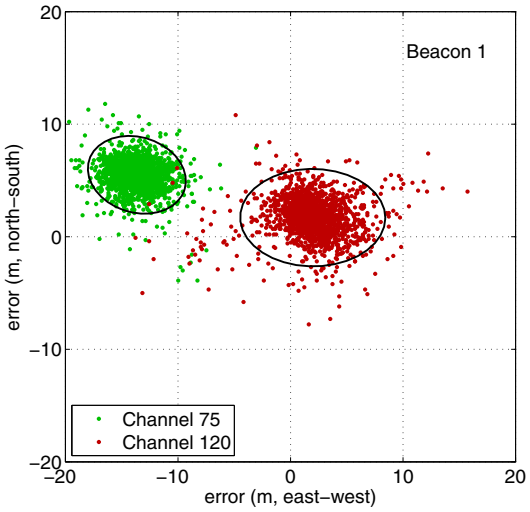

Figure 4. Localizations of a fixed-position transmitter (beacon 1) during one hour. Clock synchronization was done using two other beacon transmitters.

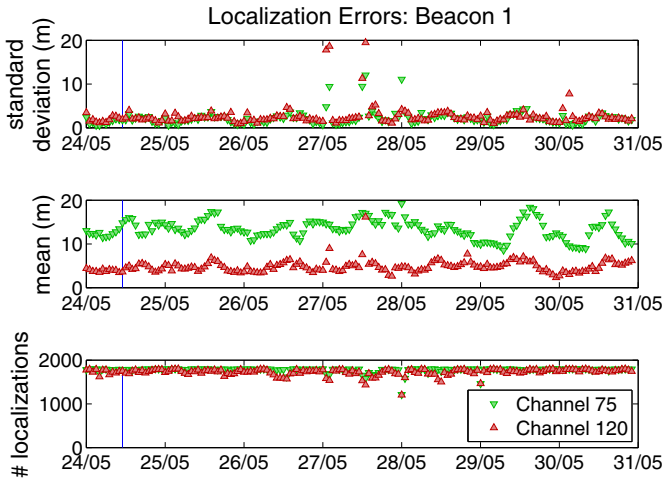

Figure 5. Localizations of Beacon 1, which transmits every second on two channels that are approximately 1MHz apart, during one particular week. The bottom graph shows the number of localizations in each hour (out of at most 1800). Times with few localizations usually indicate strong interference. The top graph shows the standard deviation of the errors (sample standard deviations over an hour) and the middle graph shows the 1-hour mean errors. The results for other beacons and times are similar. The vertical line on May 24 shows the time in which the data in Figure 4 was collected. Tick marks on the time axes indicate the beginning of each day, specified in *dd/mm* format.

tight cluster, but the clusters of the two channels are clearly separate. The centers of the clusters are several meters away from the actual location of the transmitter. Figure 5, which analyzes one week's worth of localizations from the same beacon, sheds more light on the behavior of errors. The data show that the standard deviation of the errors, estimated from one hour windows, is small (around 3m), fairly constant, and is roughly similar for both channels. The mean errors in the same one-hour windows are larger, 5–15m, and they differ significantly between the two channels, with smaller mean

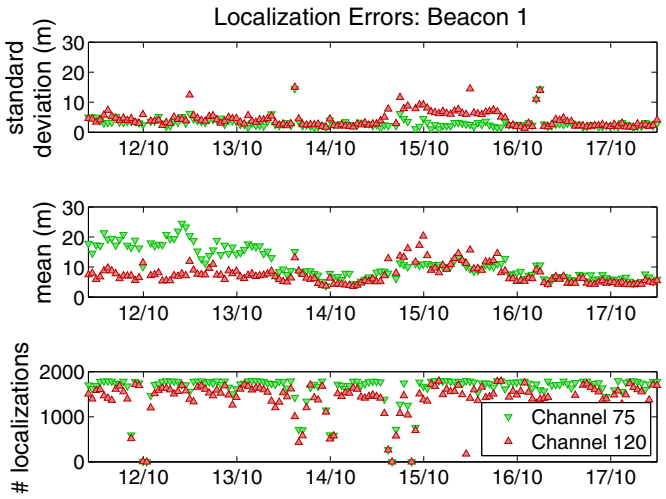

Figure 6. Localizations of Beacon 1 during another week, for comparison with Figure 5

error of around 5m in Channel 120. The data also show high system availability; most of the 1800 transmissions in each channel in each hour have been localized.

The data suggest that localization errors can be split into two main components. One component is caused by the limited information in the noisy channel and is bounded from below by the Cramer-Rao bound. This component shows little temporal correlation between localizations, it appears to be Gaussian with standard deviations of single meters, at least in the center of the coverage area. The shape of the Gaussian appears to match the confidence ellipse generated from covariance matrices of localizations (see below for more data). The other component is frequency dependent and it changes slowly.

We investigated several possible causes for the nonzero mean errors and for their frequency (channel) dependence. We believe that the slowly-varying component of the error is driven mostly by three factors: analog-delay variations in the filters of the receivers, interference, and multipath propagation, which also affects GPS and which is hard to model and eliminate [16, pages 456–457]. Figure 6 shows the statistics of the errors during another week several months later (mid October 2015). The mean errors in Channel 120 were mostly smaller than 10m, as they were in May (Figure 5). Early in the October week, errors in Channel 75 where higher than on Channel 120, as was the case in May. However, later in the October week mean errors on Channel 75 dropped, while errors in Channel 120 rose, especially during October 14 and 15. Hours with large mean errors in Channel 120 are correlated with missing localizations, which suggests that both problems are caused by strong interference on that channel. The tendency of the error to be smaller on Channel 120, which is in the middle of the band-pass of the filters in the front-end units of the receivers, leads us to believe that these filters have some contribution to the error. Channel 75 is close to the edge

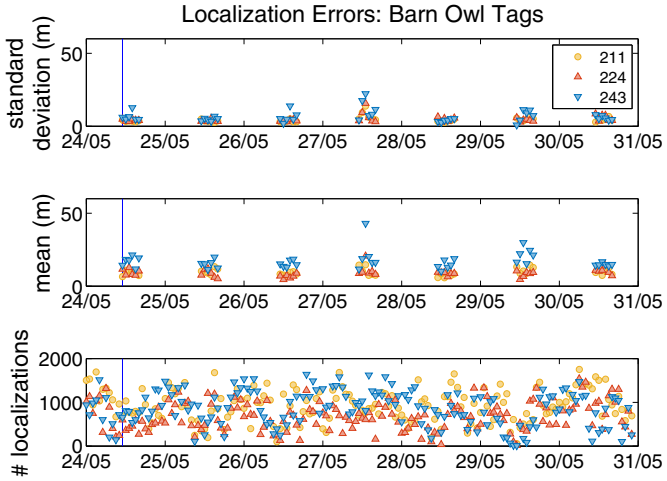

Figure 7. Localizations of tags attached to three Barn Owls during the day (11:00 to 17:00), when the birds were in their next boxes. Night-time localizations are not shown because we are not able to determine their errors. The structure of the graphs is the same as those shown in Figure 5 The tags transmitted on Channel 75 every 2s.

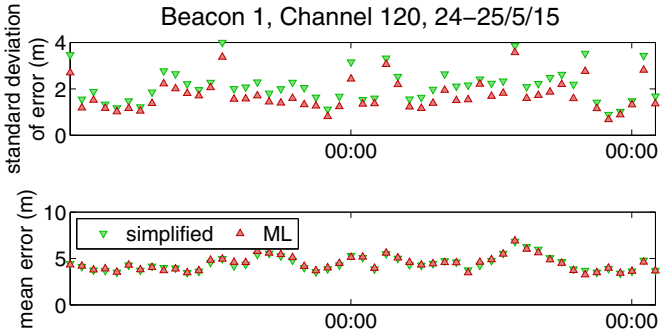

Figure 8. A comparison of localization errors with the simplified formulation and the full WLS formulation. The numbers of localizations are essentially the same in both formulations and are not shown.

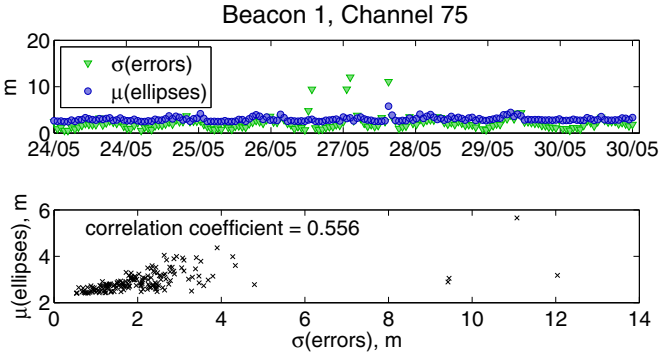

Figure 9. A comparison of the standard deviation of errors in each 1-hour window to the mean of the half-axis of 68% confidence ellipses that correspond to each localization. This data is again from the simplified algorithm.

of the band pass, where the analog delay varies more rapidly with frequency than in Channel 120.

Early on we suspected that the nonzero mean errors were caused by inaccurate specification of the location of receivers and beacons, but we dismissed this hypothesis after (1) we employed a surveyor to specify these locations and (2) we discovered that the mean error was clearly frequency (channel) dependent.

We used transmitters attached to wild Barn Owls (*Tyto alba*) to verify that ATLAS achieves similar accuracy when transmitters are attached to wild animals. We used Barn Owls because they remain relatively stationary during daylight hours, staying in a nest box at a known location. Tags attached to wild animals have lower radiated power than tags with antennas in free space, because of de-tuning of antennas and absorption of radiation by body tissue. The SNR is therefore lower than that of beacon transmissions. Also, the movement of animals (even within a nest box) causes changes in antenna position and orientation, increasing SNR variability. Figure 7 shows that the performance of ATLAS on Barn Owls is somewhat poorer than the performance on beacons, but is still very good. The standard deviation is around 5m, and the mean error is usually between 5 and 20m (the Barn Owl tags transmitted on Channel 75, which had larger mean errors even on beacons). The fraction of transmissions that have been localized is lower than with beacons, but the system still usually delivers hundreds of localizations every hour.

All the experimental results we have presented so far used the simplified localization formulation (Section II-C), which is fast enough to compute localizations at real time. The WLS formulation (Section II-B, in particular Equation 3) improves the results, especially if multiple beacon transmissions are used. Figure 8 compares the errors produced by the simplified location estimator with the errors produced by the WLS estimator. The WLS estimator used all the beacon transmissions within 30s of the transmission that is localized, and it assumes that $k_i = 0$; that is, we did not attempt to estimate clock-rate errors, which we determined to be insignificant due to the use of GPS-disciplined oscillators in receivers. Over the 51-hour span of the analysis, the WLS estimator improved the standard deviation of errors in 1-hour windows at least 78% of the time (100% for the tag and channel shown in Figure 8). The average improvement in the standard deviation ranged from 0.18m (beacon 3, channel 75) to 0.43m (beacon 1, channel 120) and was always positive. This is consistent with the theory.

The WLS estimator reduced the 1-hour mean errors by 0.99 to 1.31m on average (averaging over 51 1-hour windows) on tags attached to Barn Owls and by 0 to -0.61m on beacons (that is, on beacons the mean error grew on average).

Figure 9 shows that the correlation matrices that ATLAS computes characterize the distribution of errors in an imperfect but useful way. The figure compares the standard deviation of errors in 1-hour windows (the same data shown by the green dots in the top graph in Figure 5) to the average longer half axis of the 68% confidence ellipses derived from the

estimated covariance matrices of each localization. We can see that the two metrics are correlated, but the correlation is not perfect. For example, in the morning of May 29 both metrics increased simultaneously; also, in the afternoon of May 27 both show fairly dramatic increases. On the other hand, the large standard deviations on May 26 do not correspond to large confidence matrices (on average). This data suggests that the estimated covariance matrices, which are derived from SNRs and geometry, model a significant component of the error, but not all of it. In particular, large covariance matrices generally correspond to a high probability of large errors, but the opposite is not true; some localizations have large errors but small covariance matrices. It appears that the errors in these localizations are not caused by Gaussian noise but by some other mechanism.

V. Conclusions

We have shown that ATLAS, a distributed time-of-arrival localization system can localize transmitters attached to wild animals in a region of 10–20 kilometers in diameter. The localizations are accurate enough for many types of studies of wildlife biology and movement ecology. Errors have standard deviations on the order of 5m; the mean error in a 1-hour window is usually around 5m in the center of the bandpass of the filters in the receiver, growing to 15–24m in the edges of the bandpass. The magnitude of errors in standalone GPS receivers is similar [16].

Note that the error analysis reported here is of raw localizations (based on 4 or more receivers), without path-based filtering and smoothing. Comparisons of the performance reported here can only be made to similarly raw data of alternative systems, prior to path-based filtering and smoothing. Furthermore, we provide tools to assess data quality, which could and should be used to select particular datasets of certain estimated quality according to the intended use and research questions.

Furthermore, the covariance matrices that the system computes are good indicators of the short-term component of the error, which is caused by noise that is not correlated with the signal, but not of the slow component, which appears to be mostly due to multipath propagation (and at fringe frequencies, also by delay variations in filters). Localizations with large covariance matrices are usually inaccurate, but some localizations with small covariance matrices are also inaccurate. The system can track multiple animals simulteneously and it can produce around a thousand localizations per animal per hour. These findings are based on a real-world outdoor deployment and with lightweight tags attached to wild animals. These are significant advances over all previous non-GPS wildlife-tracking systems.

:

References

[1] Michael Allen, Lewis Girod, Ryan Newton, Samuel Madden, Daniel T. Blumstein, and Deborah Estrin. Voxnet: An interactive, rapidly-deployable acoustic monitoring platform. In *Proceedings of the 7th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 371–382, 2008.

[2] F. Cagnacci, L. Boitani, R. A. Powell, and M. S. Boyce. Animal ecology meets GPS-based radiotelemetry: A perfect storm of opportunities and challenges. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365:2157–2162, 2010.

[3] Vladimir Dyo, Stephen A. Ellwood, David W. Macdonald, Andrew Markham, Niki Trigoni, Ricklef Wohlers, Cecilia Mascolo, Bence Pásztor, Salvatore Scellato, and Kharsim Yousef. WILDSENSING: Design and deployment of a sustainable sensor network for wildlife monitoring. *ACM Transactions on Sensor Networks*, 8:29:1–29:33, 2012.

[4] K. D. Hambright and T. Zohary. Lakes Hula and Agmon: destruction and creation of wetland ecosystems in northern Israel. *Wetlands Ecology and Management*, 6:83–89, 1998.

[5] Jackson Labs Technologies Inc. FireFly 1A GPS disciplined oscillator specifications, 2010.

[6] R. Kays, S. Tilak, M. Crofoot, T. Fountain, D. Obando, A. Ortega, F. Kuemmeth, J. Mandel, G. Swenson, T. Lambert, B. Hirsch, and M. Wikelski. Tracking animal location and activity with an automated radio telemetry system in a tropical rainforest. *The Computer Journal*, 54:1931–1948, 2011.

[7] Michael A. Lombardi. The use of GPS disciplined oscillators as primary frequency standards for calibration and metrology laboratories. *Measure*, 3(3):56–65, 2008.

[8] R. MacCurdy, R. Gabrielson, E. Spaulding, A. Purgue, K. Cortopassi, and K. Fristrup. Automatic animal tracking using matched filters and time difference of arrival. *Journal of Communications*, 4(7):487–495, 2009.

[9] Robert B. MacCurdy, Richard M. Gabrielson, and Kathryn A. Cortopassi. Automated wildlife tracking. In Seyed A. (Reza) Zekavat and R. Michael Buehrer, editors, *Handbook of Position Location: Theory, Practice, and Advances*, pages 1129–1167. Wiley, 2012.

[10] Andrew Markham, Niki Trigoni, Stephen A. Ellwood, and David W. Macdonald. Revealing the hidden lives of underground animals using magneto-inductive tracking. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 281–294, 2010.

[11] Andrew C. Markham and Andrew J. Wilkinson. EcoLocate: A heterogeneous wireless network system for wildlife tracking. In Tarek Sobh, Khaled Elleithy, Ausif Mahmood, and MohammadA. Karim, editors, *Novel Algorithms and Techniques In Telecommunications, Automation and Industrial Electronics*, pages 293–298. Springer, 2008.

[12] Jorge Jesus Moré. The Levenberg-Marquardt algorithm: Implementation and theory. In *Numerical Analysis, Lecture Notes in Mathematics 630, Springer Verlag*, pages 105–116, 1977.

[13] J. A. Nelder and R. Mead. A simplex method for function minimization. *Computer Journal*, 7:308–313, 1965.

[14] M. R. Recio, R. Mathieu, P. Denys, P. Sirguey, and P.J. Seddon. Lightweight GPS-tags, one giant leap for wildlife tracking? an assessment approach. *PLoS ONE*, 6(12):e28225, 2011.

[15] Šimon Řeřucha, Tomáš Bartonička, Petr Jedlička, Martin Čížek, Ondřej Hlouša, Radek Lučan, and Ivan Horáček. The BAARA (Biological AutomAted RAdiotracking) system: A new approach in ecological field studies. *PLoS ONE*, 10:e0116785, 2015.

[16] Gilbert Strang and Kai Borre. *Linear Algebra, Geodesy, and GPS*. Wellesley-Cambridge Press, 1997.

[17] Sivan Toledo. Evaluating batteries for advanced wildlife telemetry tags. *IET Wireless Sensor Systems*, 5:235–242, 2015.

[18] Sivan Toledo. A selective robust weak-signal UHF front end. *QEX*, pages 31–36, January/February 2015.

[19] Sivan Toledo, Oren Kishon, Yotam Orchan, Yoav Bartan, Nir Sapir, Yoni Vortman, and Ran Nathan. Lightweight low-cost wildlife tracking tags using integrated transceivers. In *Proceeings of the 6th Annual European Embedded Design in Education and Research Conference (EDERC)*, pages 287–291, 2014.

[20] Don J. Torrieri. Statistical theory of passive location systems. *IEEE Transactions on Aerospace and Electronic Systems*, AES-20:183–198, 1984.

[21] Harry L. Van Trees, Kristine L. Bell, and Zhi Tian. *Detection Estimation and Modulation Theory, Part I: Detection, Estimation, and Filtering Theory*. Wiley, 2nd edition, 2013.

[22] M. Wikelski, R. W. Kays, N. J. Kasdin, K. Thorup, J. A. Smith, and G. W. Swenson. Going wild: what a global small-animal tracking system could do for experimental biologists. *Journal of Experimental Biology*, 210:181–186, 2007.