The 2010 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 18-22, 2010, Taipei, Taiwan

Autonomous Airborne Wildlife Tracking Using Radio Signal Strength

Fabian Körner, Raphael Speck, Ali Haydar Göktoğan, Salah Sukkarieh

Abstract—**In wildlife radio tagging, small radio transmitters attached to animals are located by human operators using directional antennas and analog receivers which provide audio output. The location of the transmitter is determined by listening to the signal and scanning the area while closing in. This procedure can be very tedious, especially in rough terrain. Searching radio tags with autonomous unmanned aerial vehicles (UAVs) offers a number of advantages, including better line-of-sight signal reception, terrain-independence and faster localization. In this paper we continue upon previous work by presenting a received-signal-strength (RSS) sensor implementation based on a modified commercial wildlife tracking receiver that is designed to operate on an autonomous fixed-wing UAV. Furthermore, an extension of the search and tracking framework for multiple targets that are undistinguishable from the sensors' point of view is proposed. After a brief system overview and a summary of the particle filter based approach, the signal processing theory and realization of the RSS sensor are outlined, including strategies for frequency tracking and receiver gain control. The paper also presents experimental results.**

I. Introduction

WILDLIFE radio tagging has been used by biologists for monitoring the movements of animals in behavioral studies since the late 1950s, but the basic principle has not changed significantly over time: a small transmitter being attached to the animal, the 'tag', transmits radio pulses with a few ms duration and repetition rates of 30-60 pulses per minute in the Very High Frequency (VHF) band, and the researcher uses one or more directional antenna(s) with a radio receiver to locate the tag by listening to the received signal. The search strategies used [1] are usually based on manual control of the receiver gain and antenna orientation to determine the direction of the strongest signal and triangulate the target location.

As even motorized tracking on the ground is very time-consuming, the area of coverage can be significantly increased by using airborne platforms such as light aircraft [2]. Air-to-ground tracking provides much better conditions resulting in longer range, because the fading effects in the radio wave propagation are less significant and more predictable for line-of-sight (LOS) scenarios without obstacles like vegetation, terrain, or other structures.

Furthermore, the use of small, cost-effective fixed-wing UAVs for the task reduces the overall costs of aerial tracking since existing off-the-shelf radio tag and receiver technology can be adapted to autonomous operation. However, the automatic evaluation of the tag signal from which to draw conclusions about the location of the transmitter brings about challenges in the proper modeling of radio wave propagation and their uncertainty as well as signal processing. Approaches to solving the problem with particle filters based on the observation of the RSS were proposed in [3], [4] and [5].

II. Contribution

The final goal of the research is the development of a cost-effective autonomous UAV wildlife tracking solution. This paper extends the initial work published in [3] by an approach to search and track multiple, undistinguishable targets at the same time with a relatively simple sensor implementation, and presents the theory and realization of a purpose-built RSS sensor subsystem. The next step towards an actual implementation is combining and extending common signal processing methods to a field-ready sensor. Section III briefly introduces the framework and assumptions. Section IV is concerned with a summary of particle based search and tracking and an extension for multiple targets, whereas section V introduces the RSS sensor implementation with the signal detection, receiver gain control and frequency drift compensation components. Experimental results obtained during field tests of the setup are discussed in section VI.

III. System Overview

A functional view of the UAV-based autonomous wildlife tracking system is shown in Fig. 1. From left to right there is the VHF receiver which is interfaced by a digital front-end for signal processing, the gain control unit, and the estimation of the current received tag signal strength. This value is used to update the probability density function (PDF) that represents the unknown tag location(s). As the PDF can be multi-modal and highly nonlinear, it is maintained by a particle filter (PF). The proposed guidance algorithms incorporate the PDFs to plan the UAV's trajectory on-the-fly, in order to minimize the average search time. The autopilot system carries out all lower level aircraft

Manuscript received March 10th, 2010. This work was carried out at the Australian Centre for Field Robotics. It is supported by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government.

Fabian Koerner is with the Institute of Flight Mechanics and Control, Universität Stuttgart, Germany (e-mail: fabian.koerner@tech-kyb.de).

Raphael Speck is with the Institute for Systems Optimization, Universität Karlsruhe, Germany (e-mail: raphael.speck@gmail.com).

Ali Haydar Goktogan and Salah Sukkarieh are with the Australian Centre for Field Robotics, The University of Sydney, Australia (e-mail: {agoktogan, salah}@acfr.usyd.edu.au).

control tasks. Both VHF receiver and autopilot system are proven off-the-shelf components.

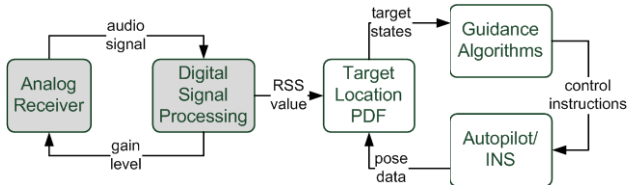

Fig. 1: Functional block diagram of the autonomous airborne wildlife tracking system.

The receiving antenna setup consists of two directional Telonics '*RA-2A'* HB9CV-antennas which are mounted horizontally below the wings of the airplane, facing towards the wingtips. The *RA-2A* provides a directive gain of 6 dBi and a front-to-back ratio of 8 dB. The feeding points are separated by half a wavelength of the VHF signal. Thus both antennas can be used at the same time with a power combiner, resulting in much higher directive gain, or switching between the left and the right side at fixed intervals with a coaxial relay, resulting in discrimination between the left and right side. The PF updates its estimated target states based on the current UAV pose and antenna orientation, the pre-calculated antenna gain pattern, an adapted '*2-Ray*' radio wave propagation model [6] and the RSS value. The fixed-wing UAV to be used is a 1:3 scale Piper J-3 Cub illustrated in Fig. 2. It is controlled by a CloudCap '*Piccolo'* autopilot system. Target localization and guidance are executed on an embedded x86 PC of the 'PC/104-Plus' form factor.

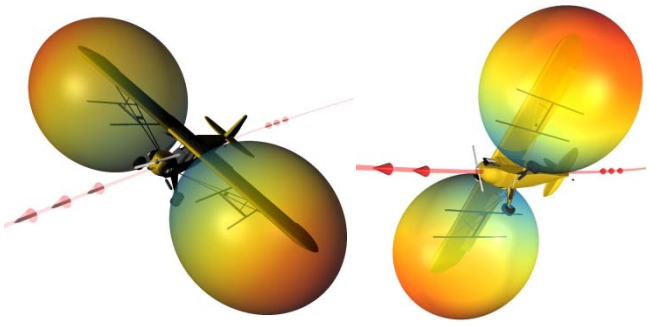

Fig. 2: 3D views of the 1:3 scale Piper J3 Cup UAV with mounted antennas and relative gain patterns. The maximum forward gain is around 6 dBi (red regions), while the wingspread is 3.2 m.

IV. Particle Filter Based Search and Tracking

A. Recursive Bayesian Estimation using a Particle Filter

A particle filter (also known as Sequential Monte Carlo filter) represents the continuous posterior PDF $p(x_t)$ by a discrete set of random samples $\chi_t = \{x_t^i, i = 0{:}N\}$ with associated normalized weights $\{w_t^i, i = 0{:}N\}$ [8, 9]:

$$p(x_t|z_{0:t}) \approx \sum_{i=1}^{N} w_t^i \cdot \delta(x_t - x_t^i) \qquad (1)$$

In this scenario, the target state x_t consists of the target position and speed. $z_{0:t}$ are the observations up to time t.

As x_t and z_t are considered first order Markov chains, the posterior PDF after a new observation depends only on the prior PDF and the observation likelihood at time t. According to Bayes' theorem, the updated PDF after a new observation is $p(x_t|z_t) = p(z_t|x_{t-1}) \cdot p(x_{t-1})/p(z_t)$, where $p(z)$ is used to normalize the posterior PDF. In the PF this update step is usually done by predicting the new state for each particle via the state transition model and updating the set of weights:

$$x_t^i = f(x_{t-1}^i, u_{t-1}^i) \qquad (2)$$

$$w_t^i = w_{t-1}^i \cdot p(z_t|x_t^i) \; \forall \; i \qquad (3)$$

The likelihood function of the observation $p(z_t|x_t^i)$ for the new particle weights is calculated by using equation (4) for the 'Detection' case and equation (5) for the 'No Detection' case, with $\Gamma_{\text{det,nodet}}$ being weighting factors and $P_{\text{RSS,expected/observed}}$ representing the RSS values from measurement and model, respectively:

$$p_{\text{det}}(z_t|x_t) \propto \Gamma_{\text{det}} \left(P_{\text{RSS,expected}} - P_{\text{RSS,observed}}\right) \qquad (4)$$

$$p_{\text{nodet}}(z_t|x_t) \propto \Gamma_{\text{nodet}} \left(P_{\text{RSS,expected}} - P_{\text{RSS,observed}}\right) \qquad (5)$$

Finally, the resampling step changes the sample distribution by duplicating particles of higher weights and dropping those of lower weights.

B. Search for Multiple Targets

In [3], a single PF was proposed for both searching and tracking of a single target. The approach can be extended for multiple targets by running several PFs in parallel. This requires the targets to be fully distinguishable, which could be achieved by using tags that either varies in pulse lengths and pulse intervals, or that sends out a unique coding of pulses [1]. This would also require a sensor implementation being able to identify the different codes.

Another option to distinguish tag signals is by using different frequencies for the tags. Since the receiver can only listen to one frequency at a time, only one PDF can be updated during that interval. Due to the very low duty cycle of a tag signal, this would make the search too slow for a realistic implementation. Thus, the ability to detect distinguishable targets is limited. But if we are more interested in the animals' distribution throughout a region or in herd behaviors, we can sacrifice the ability to clearly identify each tracked target and use tags that send on the same frequency. PFs are however poor at maintaining the multi-modality of the PDF resulting from the presence of multiple targets. Approaches have been suggested to maintain multiple modes [10], but they are not suitable for our scenario, in which observations are only available for a small part of the PDF due to the limited sensor range.

The method introduced here is able to find multiple modes in a single PDF without maintaining the multi-modality throughout the search process, but by learning from already detected objects.

For all targets whose transmitters are using the same

frequency, a single a priori PDF is used. The amount of particles can be the same as in the single target scenario, as it is still sufficient to cover the whole search area. Once a receiver detects a tag signal, the PF is split into two different filters: One PF is now used to locate the detected target (LPF), whereas the other is used to search for remaining targets (SPF).

The two filters differ in their update step: The SPF's update is solely based on negative information, i.e. it is only updated by UAVs whose sensors do not detect a signal, hence avoiding degeneration of the particle distribution. On the other hand, the PFs' behavior of degenerating quickly to a single mode even if multiple targets are within sensor range leads to the LPF detecting a single target even if multiple targets are close by or on the same spot. This means that multiple targets within sensor range get located consecutively. A target is considered to be "found" once the standard deviation of the LPF drops below a certain threshold. It gets tracked in a down-sampled PF by picking particles of the LPF with probability $p^i \propto w_t^i \; \forall \; i$.

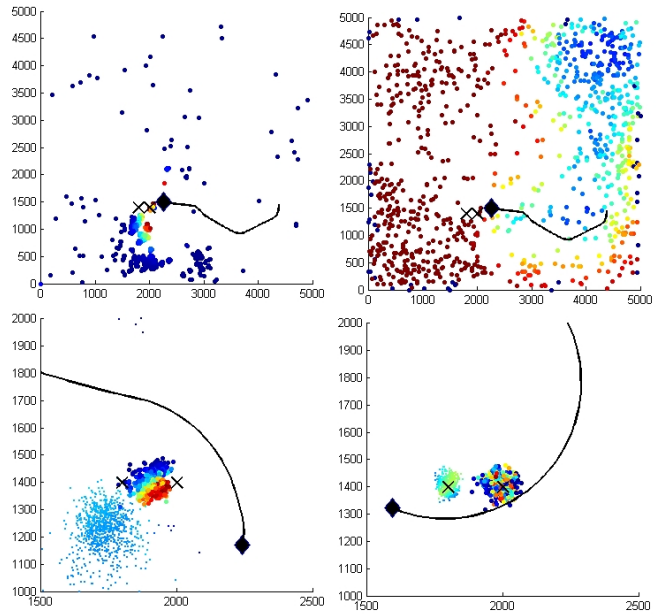

Fig. 3: Two targets with tags on the same frequency being tracked by a single UAV. Top left: LPF at t = 369 s, target detected; right: SPF at t = 369 s; bottom left: t = 506 s, one target found, second target gets located; right: t = 681 s, both targets found.

Since we now have an estimation of the target's states, the RSS expected from this target can be calculated using the PFs' internal model of the wave propagation. The RSS sensor is capable of indicating how many tag signals at a known pulse rate and length it detects, but tends to only give the RSS value of the strongest signal due to the gain control implementation described later. Consequently, the observation likelihoods for the detection case are given by equation (6):

$$p\left(z_t^j \middle| x_t^i\right) = p\left(P_{RSS,expected}^i = P_{RSS,received}^i\right) \quad (6)$$

$$P_{RSS,expected}^i = \max\left(P_{RSS,Prop.Model}^i, P_{RSS,expected}^k\right) \quad (7)$$

Here, the expected power is given by (7), and is the modeled RSS of tracked target k. The LPF also deals with the probability of any received, but not measured, signals to be above the detection threshold to accelerate degeneration of the LPF.

Fig. 3 shows the simulation results for a scenario of a single UAV searching for two targets with tags on the same frequency and a small spatial separation. The targets are represented by the black crosses, while the UAV is the black rhomboid with the line as its flight history. Plots (a) and (b) illustrate the difference between LPF and SPF, while (c) and (d) show the convergence of the two particle clouds to an accurate position for both the undistinguishable targets. This simulation serves as a proof of concept for the proposed methodology.

Algorithm 1 summarizes this method of detecting multiple non-distinguishable targets.

V. RSS Sensor Subsystem

The off-the-shelf wildlife tracking receiver being used as the primary sensor is of the type Biotrack '*Sika'*. It is a super-heterodyne receiver which down-converts a received radio signal from the VHF carrier frequency (~173 MHz in this case) to the audible range of ~0.5-4 kHz, while preserving amplitude. Standard radio tags transmit a pulsed carrier wave without further modulation, resulting in audible tones of a certain frequency within the bandwidth of the receiver's intermediate filter (*Sika*: ~3.5 kHz).

Fig. 3 shows the structure of the receiver ('ANALOG') and how it is interfaced by the digital front-end developed here ('DIGITAL'). The audio output is digitized by a 16 bit TI TLV320AIC3204 codec and processed on a fixed-point TI C5505 digital signal processor (DSP). The *Sika* gain is set by providing the control voltage for the variable gain amplifier (VGA) using a digital-to-analog converter (DAC).

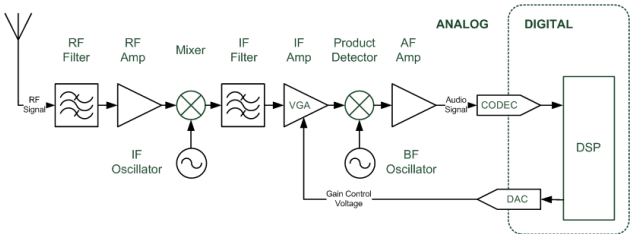

Fig. 3: Receiver structure. The gain is controlled digitally by using a DAC chip to provide the amplifier control voltage.

The update step of the particle filter requires a value of the current RSS to update the particle weights according to their location and wave propagation assumptions, as outlined earlier. This value can hardly be derived directly from the audio output as there is no continuous carrier wave, and the duty cycle of the signal is very low (e.g. ~3% for 30 ms pulses @ 60 pulses/min).

```
initialize particles for a priori PDF
while true do
  foreach tracked target
    predict particles
    update particles with all observations and predicted
      RSS from tracked targets
    predict RSS from this target
    resample if necessary
  end
  foreach SPF do
    if new target gets detected
      create LPF by cloning SPF
    else
      predict particles
      update particles with negative information only
      resample if necessary
    end
  end
  foreach LPF do
    if standard deviation > ε
      predict particles
      update particles with all observations
        and predicted RSS of tracked targets
      resample if necessary
    else
      downsample LPF to target tracking PF
    end
  end
end
```

Algorithm 1: detecting multiple non-distinguishable targets

Furthermore, the signal is prone to random fading effects, and the full gain dynamic range has to be exploited in order to achieve maximal sensitivity. The approach used here is as follows: The *Sika* gain is adjusted in a way that it follows the magnitude of the detection results, barely providing a reliable detection if a tag is in range.

In case of detection, the RSS value P_{RSS} is then derived from the sensitivity curve given in Fig. 4. The curve was measured by feeding the sensor system with a sine wave of the expected tag frequency generated by a VHF signal generator and recording the weakest input signal levels resulting in a 'True Positive' over the whole receiver gain range. In Fig. 4, the *Sika* gain G_{Sika} is represented by the control voltage of its variable gain amplifier (VGA) stage (0-3700 mV) because this value is easily accessible in the digital implementation. The nonlinearity of the curve depends directly on the characteristic curve of the VGA, and the curve therefore serves as a look-up table for the RSS estimate.

In the 'No Detection' case, the sensitivity curve defines the current sensor range: strong background noise results in a low gain and a short range, while weak background noise allows a higher gain/range without noise-induced false positives. The control loops being responsible for this behavior are explained in the following sub-section.

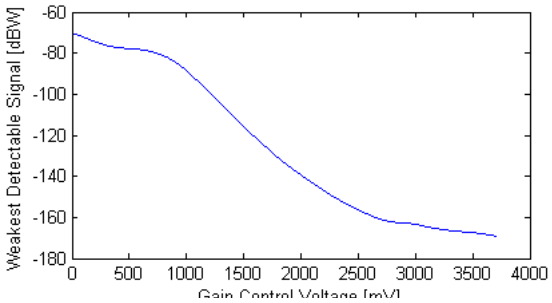

Fig. 4: Detector sensitivity curve used to obtain a coarse RSS value for the update of the PDF.

A. Receiver Gain Control

The block diagram in Fig. 5 shows the gain control approach that was implemented, taking into account the requirements of the application. It consists of two feedback loops, one for the receiver gain control and one for a digital factor applying on the reference signal of the other loop.

The receiver gain control loop incorporates the *Sika* VGA, (represented by a multiplier), a low-pass filter and the 'Receiver Gain Calculation' block. After digitization, the audio signal is low-pass filtered with a cutoff frequency well below the expected tone frequency. The filtered signal is then used as reference for the gain calculation, whose main purpose is to exploit the full dynamic range of the audio codec without triggering the detector with strongly amplified background noise in absence of a tag signal.

The digital gain factor that is controlled by the outer loop applies on the reference for the receiver gain loop according to the current detection result. In case of detection, it steps in by increasing the reference level and therefore forcing the gain down. The strategy aims on keeping the detection filter output within a certain range (e.g. between 100 % and 150 % of the detection threshold), so that the adjustment of the receiver gain follows the peaks caused by the tag pulses rather than the background noise level. This behavior prevents strong signals at close range from clipping and allows the approximation of the current RSS value from the gain level, as described above. In the 'no detection' case, the digital gain factor is decreased slowly over several pulse repetition periods with a constant rate.

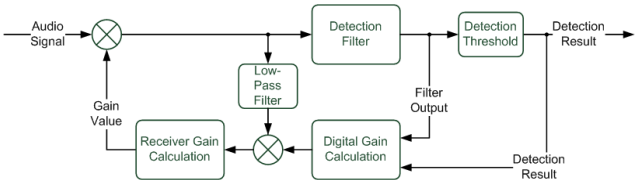

Fig. 5: Receiver gain control loops. The receiver gain loop controls the signal level for the digitization, while the digital gain control steps in to follow the strength of detected tones.

B. Signal Detection

The 'Goertzel' algorithm [11] is used for the detection of the tag tones within the audio signal, which is essentially an efficient means to compute spectral components ('bins') of a

signal without having to run a whole discrete Fourier transformation (DFT). It takes advantage of the periodicity of the phase factors, which allows expressing the computation of the DFT as a linear filter operation. Representing the resulting complex one-pole filter by a conjugate complex two-pole filter with only real multiplications results in the difference equation (8) for each filter step and new sample x_n:

$$v_n = x_n + a\, v_{n-1} - v_{n-2} \quad (8)$$

$$y = v_{n-1}^2 + v_{n-2}^2 + v_{n-1} v_{n-2} - a \quad (9)$$

The filter coefficient $a = 2\cos(2\pi f / f_s)$ arises from the frequency of the tone f and the sampling frequency f_s. The result of the algorithm, y, which represents the power of the bin, is computed after N iterations, where N is the overall filter length. As no phase information is required, the originally complex calculation of y can be simplified to few real operations given by equation (9). A block diagram of the Goertzel filter is given in Fig. 6, where the left part is an infinite impulse response (IIR) filter representation of (8) and the right part within the dotted box represents (9), being executed every N steps only. The bandwidth and frequency resolution of the filter is given by $B = f_s/N$.

The computational effort for real signals is N multiplications and $2N$ additions each filter step plus additional 4 multiplications and 3 additions overhead for the calculation of y after N steps. In case of a variable tone frequency, the filter coefficient a has to be recalculated each time the search frequency is updated as well.

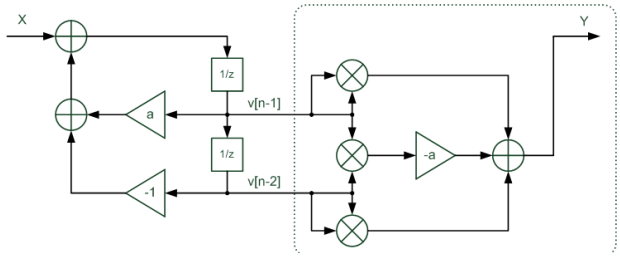

Fig. 6: Block diagram of the Goertzel filter for selective spectral analysis.

C. Frequency Drift Compensation

Most radio tags are simple analog circuits which are optimized for small size, low weight and low power consumption. Therefore they are prone to frequency drifts caused by temperature. Changes in the transmission frequency of the tag result in a change of the tone at the receiver audio output. As the detector algorithm necessarily has a narrow frequency response, the search frequency has to be automatically adapted to the changing conditions. The approach is as follows: Three detection filters tuned to slightly different frequencies are executed in parallel and their results are compared. The first filter is called 'Low' and searches for a slightly lower frequency $f_L = f - \Delta_f$ than the second filter, 'Proper', which is tuned to the center frequency $f_P = f$. The third filter, 'High', searches for the slightly higher frequency $f_H = f + \Delta_f$. The left part of Fig. 7 illustrates the frequency responses of the resulting filter setup schematically.

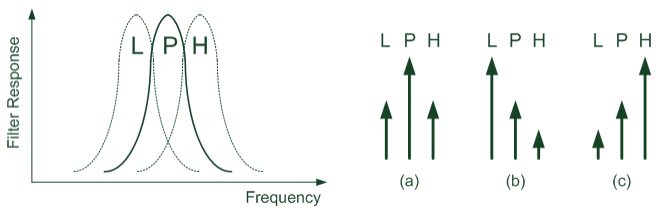

Fig. 7: Frequency tracking. Left: frequency responses of the three parallel filters, right: three possible results in case of a detection.

Three cases can occur during operation, resulting in a filter bank that converges to the correct frequency:

(a) *P maximal*: no frequency change
(b) *L maximal*: all three frequencies are decreased
(c) *H maximal*: all three frequencies are increased

The frequency separation Δ_f and the frequency increment or decrement step have to be chosen carefully according to the filter bandwidth, e.g. $\Delta_f = 0.5\,B$ and $f_{step} \ll B$, to prevent the system from oscillation. An initial guess for the current transmission frequency of a tag that is deployed in the field can be made from the ambient temperature, assuming that the frequency drift over temperature is known from prior measurements. If the guess is as close as $1.5\,B$ to the real value of f, the system will then converge to the proper center frequency resulting in optimal signal detection performance.

VI. Experimental Results

To validate the RSS sensor implementation, ground based field tests were performed. The antenna setup was mounted on the rooftop of a car in the same manner as it would be on the UAV, a Biotrack '*TW-3*'-type radio tag was placed at different locations, and the sensor output was recorded with 2 Hz while moving through the area. The receiver gain, active antenna, detection result, search frequency, position, speed and orientation were logged.

Fig. 8 illustrates the signal reception performance for three target locations. Each plot contains the superposed results of several rounds. The arrows indicate the orientation of the receiving antenna(s), while the color gives the detector output: green for 'Detection', red for 'No Detection'. The arrow length is inversely proportional to the gain setting and indicates the gain level / RSS, where green and red arrow lengths are normalized independently to highlight the detections. It can be seen that the background noise is stronger in the upper left corner, resulting in lower gain. Generally, the ground-to-ground range is poor, and signal detection is only possible in few cases due to the directional pattern of the antennae. However, there are no false positives, a fact that significantly improves the performance of the PF as it was introduced earlier. The gain control and signal detection approaches are therefore

generally valid.

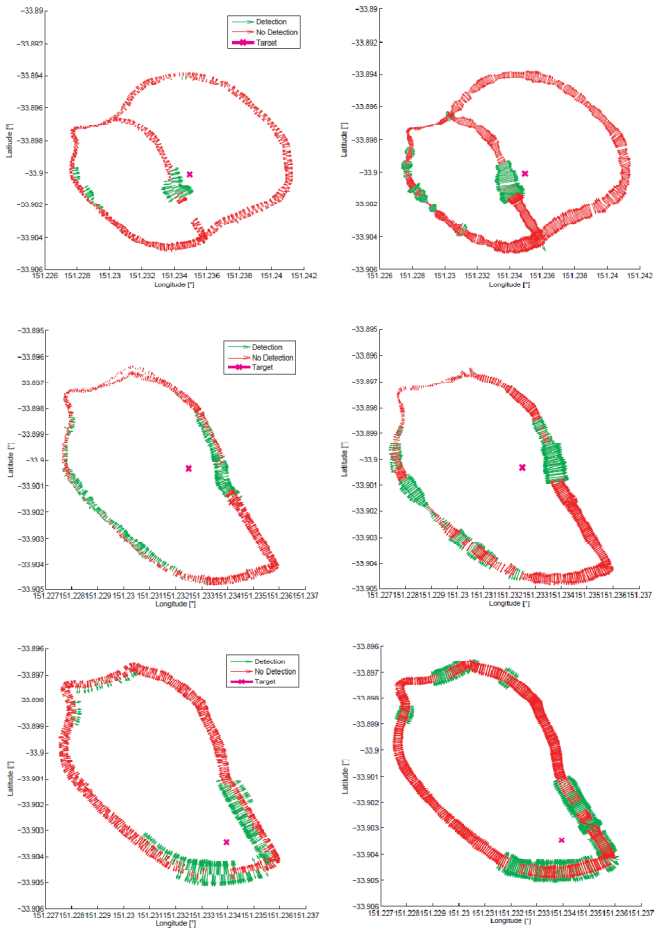

Fig. 8: Signal reception for three different target locations. Left: antenna toggling, right: both antennas active simultaneously.

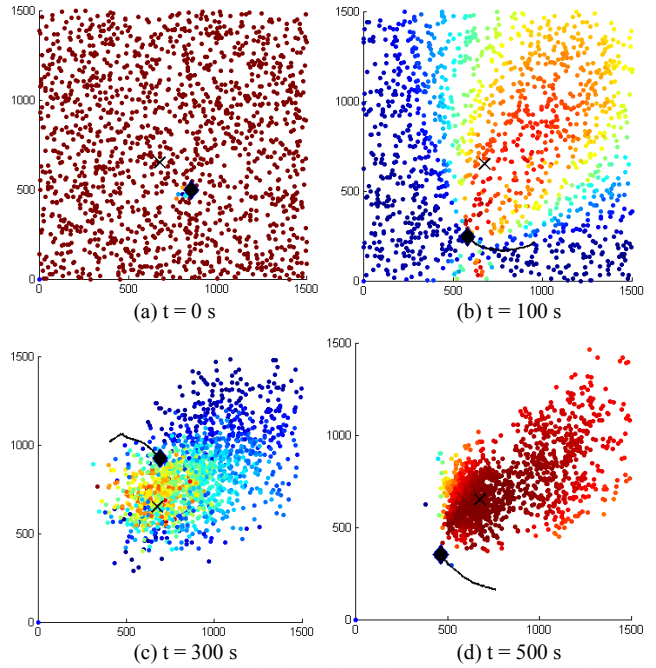

Fig. 9: The particle cloud calculated from the recorded field test data for different time steps. The x/y axes are in m.

The recorded data set shown in the second row of Fig. 8 was post-processed by using the implemented particle filter. Fig. 9 shows the particle representation of the estimated target position for the dataset at different time steps. Plot (a) shows the uniform distribution at $t = 0$, (b) illustrates the massive change in the particle weights according to the antenna gain pattern on the first signal detection, while (c) and (d) show the convergence of the cloud to the real target position (x) due to multiple observations from several directions.

VII. Conclusions

Extensions of the theoretical search and tracking framework based on the initial work published in [3] were presented focusing on multiple targets and were validated by simulation. Furthermore, a RSS sensor system for the given scenario was developed and evaluated during field tests. The results show that the approach is generally applicable, but there is still potential for optimization. Applying the system on an airborne scenario in a rural environment will result in higher range due to LOS conditions and weaker background noise. The use of more sophisticated radio tags with a higher duty cycle and/or coded signals designed to be detected under low signal-to-noise conditions would improve the detection performance and range as well.

References

[1] Kenward, R., A Manual for Wildlife Radio Tagging, London: Academic Press, 2001.

[2] Seddon & Maloney, 'Tracking wildlife radio-tag signals by light fixed-wing aircraft', Technical report, Wellington/New Zealand: Department of Conservation, 2004.

[3] Posch, A. & Sukkarieh, S., 'UAV based search for a radio tagged animal using particle filters', Australasian Conference on Robotics and Automation (ACRA), December 2-4, 2009, Sydney, Australia, 2009.

[4] Soriano; Caballero & Ollero, 'RF-based Particle Filter localization for Wildlife Tracking by using an UAV', 4th International Symposium on Robotics, 2005.

[5] Caballero; Merino; Maza & Ollero, 'A Particle Filtering method for Wireless Sensor Network Localization with an Aerial Robot Beacon', IEEE International Conference on Robotics and Automation, pp. 596 - 601, 2008.

[6] Blaunstein, N. & Christodoulou, C., 'Radio Propagation and Adaptive Antennas for Wireless Communication Links', Wiley Interscience, 2007.

[7] Thrun, S.; Burgard, W. & Fox, D., 'Probabilistic Robotics', The MIT Press, 2005.

[8] Arulampalam et al., 'A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking', IEEE Transactions on Signal Processing 50(2), p. 174-188, 2002

[9] Chung, C. F. & Furukawa, T., 'Coordinated pursuer control using particle filters for autonomous search-and-capture', Robotics and Autonomous Systems (2009), 2009.

[10] Vermaak, J.; Doucet, A. & Perez, P., 'Maintaining Multi-Modality through Mixture Tracking', Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003), 2003.

[11] Schmer, M., 'DTMF Tone Generation and Detection: An Implementation Using the TMS320C54x', Technical report, Texas Instruments, 2008.