

This item is the archived peer-reviewed author-version of:

A biomimetic radar system for autonomous navigation

Reference:

Schouten Girmi, Steckel Jan.- A biomimetic radar system for autonomous navigation
IEEE transactions on robotics - ISSN 1552-3098 - (2019), p. 1-10
Full text (Publisher's DOI): <https://doi.org/10.1109/TRO.2018.2889577>

A Biomimetic Radar System for Autonomous Navigation

Girmi Schouten, *Graduate Student Member, IEEE*, Jan Steckel, *Member, IEEE*

Abstract—This paper presents a novel biomimetic radar sensor for autonomous navigation. To accomplish this, we have drawn inspiration from the sensory mechanisms present in an echolocating mammal, the common big-eared bat (*Micronycteris microtis*). We demonstrate the correspondence in both the hardware, system model and signal processing. To validate the performance of the sensor we have developed a complementary control system based on subsumption architecture, which allows the system to autonomously navigate unknown environments. This architecture consists of separate behaviors with different levels of complexity, which are combined to produce the overall functionality of the system. We describe each behavior separately and examine their performance in real-world navigation experiments. For this purpose, the system is placed in two distinct office environments with the goal of achieving smooth and stable trajectories. Here we can observe noticeable improvements when employing high-level behaviors. Furthermore, we utilize the data collected during the navigation experiments to perform simultaneous localization and mapping, using an algorithm developed in our earlier work. These results show a substantial improvement over the odometry. We attribute this to the fact that the system traverses stable and repetitive paths, which facilitates place recognition.

Index Terms—Radar Sensing, Biologically-Inspired Robots, Behavior-Based Systems, Autonomous Agents, SLAM

I. INTRODUCTION

AS automated control of vehicles is moving towards full autonomy, improvements in both hardware and software are needed to allow these systems to operate more robustly in a wide variety of environments and weather conditions [1]. For this purpose autonomous vehicles can employ a broad range of sensors such as lidar, radar, and cameras. Each of these sensing modalities has its own merits and shortcomings, and none offers reliable performance under all possible conditions. Therefore, a combination of multiple types of sensors is needed to safely and robustly sense the environment [2]. Furthermore, each system requires its sensors to adhere to specific constraints, such as size, range, or price. This further increases the need for a diverse assortment of sensors from which a subset can be selected to suit the requirements at hand.

In this paper we present the first iteration of a biologically-inspired radar sensor for autonomous navigation. This sensor aims to deliver a compact, low-cost solution for remote sensing under challenging environmental conditions. To achieve the

Girmi Schouten is with the CoSys-Lab at the Faculty of Applied Engineering, University of Antwerp, Belgium, and with the Flanders Make Strategic Research Centre. He is also an SB PhD fellow at FWO (Research Foundation - Flanders), project [1S09417N]. (email: girmi.schouten@uantwerpen.be)

Jan Steckel is with the CoSys-Lab at the Faculty of Applied Engineering, University of Antwerp, Belgium, and with the Flanders Make Strategic Research Centre. (email: jan.steckel@uantwerpen.be)




Fig. 1: Comparison between biological sonar and biomimetic radar sensing. (a) A common big-eared bat (*Micronycteris microtis*) with its ears (orange) and mouth (green) marked. Image courtesy of Inga Geipel, Smithsonian Tropical Research Institute, Panama. (b) The radar sensor used in this work. Analogous to the bat, the receiver and emitter antennas are marked using the same color coding.

aforementioned objectives, the sensor uses radar technology and is constructed using a minimal amount of components to reduce its overall form-factor as well as inherent manufacturing cost. The inspiration for this approach is taken from nature, where bats show that comprehensive perception of the environment is possible, using only their ears and mouth as interface to the external world. Not only can bats reliably find their way back and forth between roosting and feeding grounds, they can also catch insect prey on the wing, and even discern between multiple types of flowering plants [3], [4]. In the past our research group has already proven the capabilities of a biomimetic sonar system both for navigation [5] as well as simultaneous localization and mapping (SLAM) [6]. Now we apply our knowledge on the acoustic domain to the electromagnetic domain. Both the hardware and software of the system are based on the main principles of biomimetic sonar; the design of the sensor resembles the morphology of an echolocating bat, while the signal processing focuses on echo features such as time-of-arrival and spectral content. An illustration of the physical similarity between both sensing methods is given in Figure 1. The reason for the shift from sonar to radar is because the latter offers some important benefits over the former; radar performs better in the presence of environmental hazards such as rain, fog, and smoke [7]. It also suffers less from atmospheric attenuation as well as other influences such as air temperature and humidity [8]. Additionally, radar signals travel at the speed of light (2.99×10^8 m/s), while sonar signals travel at the speed of sound (3.43×10^2 m/s). This means echoes return faster, which in turn allows the system to be more responsive.

Complementary to the biomimetic radar sensor, we have also developed a control architecture to demonstrate the system's capabilities. This takes the form of a subsumption architecture [9], which uses the radar sensor as the sole input

device to achieve Sense-and-Avoid (SAA) functionality. It is a type of reactive control architecture which can be situated in the field of behavior-based robotics [10]. In this case, inspiration is taken from insect intelligence, which proves to be successful in real-world navigation despite its seemingly low complexity.

Lastly, the sensory data gathered during autonomous navigation is used as input for a graph-based SLAM solution which we developed in previous work [11]. This solution is based on the navigational processes present in the mammalian hippocampus, and uses as input a spectrogram representation of the radar signals, similar to the information presumed to be available to bats. Using a separate sensor setup to gather ground truth information, consisting of a lidar sensor and existing SLAM algorithm, we generate a metric map of the environment and determine the exact path of the robot within the map. This information is then employed to assess the results of our experiments.

To summarize, the novelty of this work with respect to previous realizations lies in the design of a single channel radar system with biomimetic properties in both hardware and software, and its application to autonomous navigation using a newly-developed control architecture. The rest of the paper is structured as follows; Section II details the workings of the sensor and the processing of the radar signals, Section III describes the design and implementation of the control architecture, Section IV discusses the results of the autonomous navigation and SLAM experiments, and Section V draws a conclusion on the system as a whole and its real-world performance.

II. RADAR SENSING

For the basis of our work, we use Flat Earth’s Ancho Radar Development Kit [12]. It combines Novelda’s XeThru X2 radar SoC, a BeagleBone Black, and sinuous antennas [13]–[15]. The sensor is capable of emitting signals using one of 10 configurable pulse generators with center frequencies between 4.4 and 10.3 GHz and respective bandwidths from 1.7 to 3.1 GHz. These large bandwidths allow for the emission of very short pulses of around 1.5 ns, which corresponds to a theoretical range accuracy of approximately 4 mm. Furthermore, the sensor has a sampling rate of 39 GS/s, which allows it to capture the full waveform of incoming electromagnetic signals. Finally, a Cilantro switching cape demultiplexes the single receiver and transmitter channels of the Ancho kit to two receivers and transmitter channels each, resulting in a total of four selectable transceiver pairs [16].

The hardware arrangement of our sensor is inspired by the morphology of an echolocating bat, as shown in Figure 1; bats use their mouth or nose as an emitter and their two ears as receivers for ultrasonic signals. In our system, a single transmitter antenna and two receiver antennas are used. For each measurement, the system sends out an electromagnetic pulse through the center antenna and records echoes through one of the antennas on the side. The Cilantro cape is then switched to the receiving antenna on the opposite side and the process is repeated. The receiving antennas are pointed slightly sideways

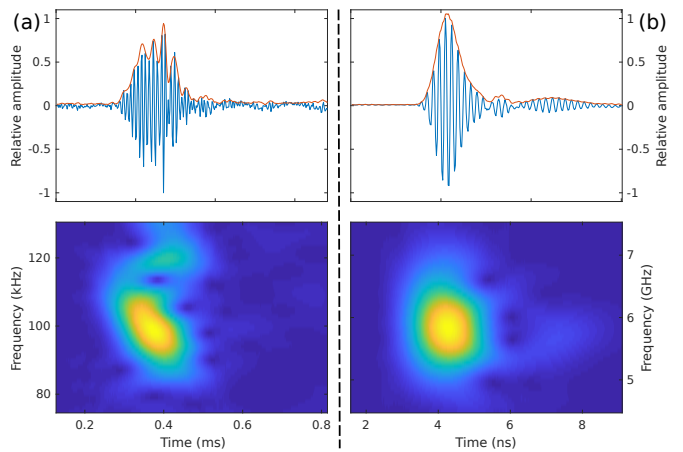


Fig. 2: Recorded echo signals (top, blue), signal envelopes (top, orange) and corresponding spectrograms (bottom). (a) Call made by a common big-eared bat (*M. microtis*). (b) Pulse generated by the radar sensor.

(30°), similarly to the pinnae of a bat. Because of this layout, it is possible to derive positional information of reflectors in the horizontal plane, using interaural time difference (ITD) and interaural intensity difference (IID), which are key features in biological sound localization [17].

A. Radar system model

The radar emits a signal $s_e(t)$, which is reflected by objects in the environment and picked up by the receivers. For the Ancho radar with the X2 SoC, $s_e(t)$ is a Gaussian-modulated sinusoid:

$$s_e(t) = s_g(t) \cdot \sin(2\pi \cdot f_c \cdot t)$$

in which f_c is the central frequency of the pulse generator (6 GHz), and $s_g(t)$ a Gaussian window function:

$$s_g(t) = \exp\left(-\frac{(t - \Delta t)^2}{2\sigma^2}\right)$$

with Δt the time shift of the window function, equal to 0.75 ns, and σ the width of the window, equal to 0.25 ns.

Fitting the biomimetic approach, this type of signal corresponds to that of the Egyptian fruit bat (*Rousettus aegyptiacus*) [18]. However, to maintain consistency throughout the paper, we compare it to the call of the *Micronycteris microtis* in Figure 2. Although this call is in actuality a very short frequency-modulated sinusoid, the signals still have a strong resemblance; both occupy a wide bandwidth and have a narrow peak in their autocorrelation function. This similarity extends to the spectrograms, where both signals cover a large frequency range and short timespan.

The radiation patterns of the emitting and receiving antennas also play an important role in localization, because they alter the outgoing and incoming signal depending on the frequency and angle of incidence. Because we are dealing with an active sensing mechanism, the effects of both the emission and reception channel are present in each echo. Therefore, we will be looking at them as a single phenomenon, which is described by the echo-related transfer function (ERTF) [19]. It is based on the head-related transfer function (HRTF), which describes the linear time-invariant filtering a signal undergoes when it

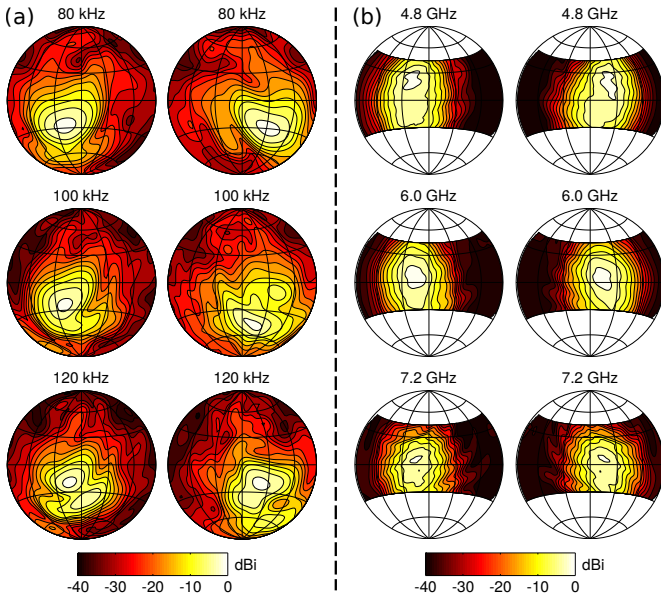


Fig. 3: Comparison between the echo-related transfer function (ERTF) of a bat and the radar sensor. A Lambert azimuthal equal-area projection is used with gridlines spaced at 30° . Plots are normalized to 0 dB and contour lines are spaced 3 dB apart. (a) ERTF of a common big-eared bat (*M. microtis*) at selected frequencies. Columns correspond to the left and right ear respectively. (b) ERTF of the radar sensor at selected frequencies. Columns correspond to the left and right receiver antenna respectively.

is picked up by a receiver. The ERTF extends this notion by taking into account the combined filtering of both the receiver and emitter in a monostatic system. This process is defined as follows; first, the emitted pulse is filtered by the frequency dependent radiation pattern of the emitter antenna for every direction ψ , modeled using a time-domain impulse response $h_e(t, \psi)$. Next, the emitted signal is reflected by objects in the environment, adding an unknown reflector filtering $h_r(t)$. Upon reception the signal is filtered by the left and right radiation patterns, again modeled using impulse responses $h_L(t, \psi)$ and $h_R(t, \psi)$. The final received signals by the left and right radar channels can be modeled as a superposition of the individual signals due to the linearity of the signal propagation model:

$$s_L(t) = \sum_{n=1}^N h_L(t, \psi_n) * h_r(t) * h_e(t, \psi_n) * s_e(t - \Delta t(n))$$

$$s_R(t) = \sum_{n=1}^N h_R(t, \psi_n) * h_r(t) * h_e(t, \psi_n) * s_e(t - \Delta t(n))$$

with $*$ denoting the time-domain convolution and $\Delta t(n)$ the time-delay caused by the two-way travel time of the emitted signals for the n^{th} reflector:

$$\Delta t(n) = \frac{2 \cdot d_n}{c}$$

with c the speed of light and d_n the distance between the radar sensor and the n^{th} reflector. Using a frequency domain model of this process, the multiplicative effect of the individual

directivity patterns on the received signals becomes more clear:

$$H_L^E\{j\omega, \psi_n\} = H_e\{j\omega, \psi_n\} \cdot H_L\{j\omega, \psi_n\}$$

$$H_R^E\{j\omega, \psi_n\} = H_e\{j\omega, \psi_n\} \cdot H_R\{j\omega, \psi_n\}$$

where H_L^E and H_R^E are the echo-related transfer functions of the left and right reception channels. Note that uppercase letters indicate the Fourier-transformed counterparts to their time-domain form. Next, these are multiplied with the spectrum of the emitted pulse and the transfer function of the reflector to arrive at the received signal spectrum at the left and right antenna, represented by S_L and S_R :

$$S_L\{j\omega, \psi_n\} = H_L^E\{j\omega, \psi_n\} \cdot S_e\{j\omega\} \cdot H_r\{j\omega\}$$

$$S_R\{j\omega, \psi_n\} = H_R^E\{j\omega, \psi_n\} \cdot S_e\{j\omega\} \cdot H_r\{j\omega\}$$

This means that it is possible to correlate the spectrum of the received signal S to the direction ψ of the originating reflector, if the spectrum of the emitted signal S_e , the ERTF of the system H^E and the received signal S are known. This holds under the assumption that a reflector does not apply a significant amount of spectral distortion. The specific set of changes induced by the ERTF in a received signal are referred to as spectral cues. The ERTF of our radar system, seen in Figure 3b, was gathered by mounting the sensor on a pan-tilt unit opposed to a radar retroreflector in an electromagnetic anechoic chamber. By analyzing the spectral content of the reflected signals for different orientations of the sensor, the ERTF is determined. Figure 3a shows the ERTF of a bat, which was computed using the Boundary Element Method (BEM) and the data presented in [20]. When comparing these ERTFs, we can observe the following; the direction of the main beam shifts slightly with respect to frequency in both cases. The radar also exhibits a similar beam width, but with a much steeper drop-off compared to the bat. Additionally, due to the antennas being frequency independent, their ERTF is significantly more stable over a large range of frequencies. This lack of variability is actually a disadvantage because less information can be extracted from the spectral content of the signal. Because our system currently only navigates in 2D, this is not crucial to its functionality. However, when performing SLAM this information can be valuable as the specific filtering by distinct reflectors enables the system to recognize particular objects and structures, which facilitates localization.

B. Signal processing

Two main signal processing techniques are applied to the radar signals to produce the information used by the control architecture. The first is an envelope function, $s^{env}[n]$. More specifically, it is the amplitude of the analytic signal generated using the discrete Hilbert transform \mathcal{H} [21]:

$$s_L^{env}[n] = \sqrt{s_L[n]^2 + \mathcal{H}(s_L[n])^2}$$

$$s_R^{env}[n] = \sqrt{s_R[n]^2 + \mathcal{H}(s_R[n])^2}$$

with $s_L[n]$ and $s_R[n]$ being the signals sampled at the left and right receiver respectively, at time step n . Peaks in this signal

correspond to received echoes of objects in the environment. By taking into account the position of these peaks, the sampling rate of the sensor and the speed of light, it is possible to determine the distance of a reflector to the sensor. Additionally, by looking at the difference in time and amplitude between peaks in the signals at the left and right receiver, the direction of a reflector can also be inferred. These properties encode information on obstacles in the environment and therefore serve as viable inputs for sense-and-avoid functionality which enables a robot to autonomously navigate an unknown area. In our system they are used by several behaviors in calculating the behavior's activation conditions and velocity outputs.

The second processing technique used is the short-time Fourier transform (STFT) [22]. It estimates the time-frequency distribution of the signal, similar to the filter bank functionality of the mammalian cochlea [23].

$$\mathcal{X}_L[m, k] = \sum_{n=-\infty}^{\infty} s_L[n] \cdot w[n - m] \cdot e^{-j2\pi kn/N}$$

$$\mathcal{X}_R[m, k] = \sum_{n=-\infty}^{\infty} s_R[n] \cdot w[n - m] \cdot e^{-j2\pi kn/N}$$

Here $\mathcal{X}[m, k]$ is the STFT at time step m and discrete frequency k . In our case, $w[n]$ is a Hamming window with a length N of 128 samples. The overlap between consecutive windows is 127 samples or Δm equal to 1 sample. A requirement to perform (short-time) Fourier analysis is that the signal under consideration is sampled above its Nyquist frequency. Because the radar sensor has a sampling rate of 39 GS/s, while the signal contains maximum frequencies of 10 GHz, the electromagnetic waveform can be reconstructed unambiguously, and thus the system fulfills the condition. For our purpose, we take the magnitude of the complex-valued STFT:

$$\mathcal{S}_L[m, k] = |\mathcal{X}_L[m, k]|$$

$$\mathcal{S}_R[m, k] = |\mathcal{X}_R[m, k]|$$

resulting in a spectrogram \mathcal{S} , which represents the estimated intensity of each frequency at each time step.

A spectrogram can be used to observe the spectral content of a signal while retaining temporal information. In this work we use spectrograms to properly calculate the change in signals over time. This is necessary because subsequent signals can vary in both time and frequency. For example, when moving toward a reflector the received echo changes in its time of arrival. On the other hand, the reflection of a corner changes in spectral content depending on its relative orientation with respect to the sensor. Both these phenomena are simultaneously captured in spectrograms.

III. CONTROL ARCHITECTURE

Subsumption architecture is a behavior-based control architecture, designed in a bottom-up, layered manner. Most classical control systems follow a sequence of operations which are executed in a fixed order. This usually consists of the following steps; sensing the external world, building and updating the internal model of that world, deciding on




Fig. 4: Schematic representation of the subsumption architecture. Information flows from left to right; a radar measurement is performed and processed for both channels. This information is passed on to each of the behaviors, which signal their intent to the arbitrator. A single behavior is then selected to take control of the robot, based on priority.

which action to take, taking into account the effects of that action, and executing the actual action. On the other hand, a system built according to subsumption architecture is made up of layers of separate behavior which execute in parallel. These behaviors consist of tightly-coupled sensorimotor loops. Each behavior has its own priority and tries to achieve its own goal. Behaviors are given control according to their priority; low-level behaviors, such as stopping when an object is near, take precedence over high-level behaviors such as reaching a specific location. Only when there is no object near, does the low-level behavior relinquish control and can the high-level behavior take over. Although separate behaviors do not directly communicate with one another, they can sense each other's actions through the external world and will react accordingly. For our specific implementation of subsumption architecture, an arbitration scheme is used. In this scheme, the arbitrator selects a single behavior to take control of the mobile robot, based on each behavior's state and priority. An overview of the implementation of subsumption architecture in our system can be seen in Figure 4. At each time step, the radar performs a measurement for both the left and right channel of the sensor. The system calculates the envelope and spectrogram of these signals according to the methods described in Section II-B. These signals are then passed on to the appropriate behaviors and each behavior decides if it needs to activate under the current conditions. The arbitrator then selects the active behavior with the highest priority in the system and collects its intended velocity vectors, which are passed on to the robot. This architectural layout is inspired by previous work of Steckel and Peremans regarding robot control using acoustic flow [24]. It should be noted that similarities restrict themselves solely to the layout, as the previously mentioned work uses a sonar imaging sensor which enables explicit 3D localization of reflectors, while the current work uses a biomimetic radar sensor which implicitly localizes reflectors through the use of temporal and spectral cues. As such the sensorimotor control laws underlying these architectures differ greatly.

The rest of this section will specify the individual behaviors in order of priority. Each section has a mathematical notation of the rules governing the velocity outputs. The variables used in these equations are detailed in Table I. Here, the distance

TABLE I: Behavior variables

Variable	Explanation
V_l	Linear velocity produced by the behavior
V_a	Angular velocity produced by the behavior
V^{min}, V^{max}	Minimum and maximum value for a velocity component
d_r	Distance to closest reflector
Δr	Difference in distance between closest reflector at left and right channel
Δs	Difference in temporal-spectral change between signals at left and right channel
T_A	Threshold of the amplitude for which the behavior activates, defined experimentally
T_d	Threshold of the distance for which the behavior activates
C	Activation condition of the behavior, as binary value

to the closest reflector is determined by:

$$\begin{aligned} d_L &= \min(n) : n < T_d \wedge s_L^{env}[n] > T_A \\ d_R &= \min(n) : n < T_d \wedge s_R^{env}[n] > T_A \\ d_r &= \min(d_L, d_R) \end{aligned}$$

with d_L and d_R the distance to the first peak above threshold T_A within distance T_d in the envelope s^{env} of the signal at the left and right receiver respectively. The difference between the closest reflectors is determined by:

$$\Delta r = d_L - d_R$$

These equations imply that the system responds the closest reflector encountered at each channel, which does not require matching of multiple, distinct echo signals.

The difference in temporal-spectral change is determined by looking at difference in change of subsequent spectrograms between both channels:

$$\begin{aligned} \Delta s &= \sum_{m=1}^M \sum_{k=1}^K |S_R^t[m, k] - S_R^{t-1}[m, k]| - \\ &\quad \sum_{m=1}^M \sum_{k=1}^K |S_L^t[m, k] - S_L^{t-1}[m, k]| \end{aligned}$$

with $S^t[m, k]$ the spectrogram of the signal at the current time step t , with size $M \times K$.

A. Collision detection (CD)

Collision detection is the most low-level behavior in the system and thus has the highest priority. Its goal is to detect and handle situations where a collision is (almost) imminent. It does this by applying a backwards throttle to prevent or minimize the impact of a collision when an object is detected at close range, with the amount of throttle depending on how close the object is:

$$\begin{aligned} C &= (n < T_d \wedge (s_L^{env}[n] > T_A \vee s_R^{env}[n] > T_A)) \\ V_l &= V_l^{min} + (V_l^{max} - V_l^{min}) \cdot \frac{T_d - d_r}{T_d} \\ V_a &= 0 \end{aligned}$$

with T_d corresponding to 0.1 m, V_l^{min} equal to -0.1 m/s and V_l^{max} equal to -0.8 m/s.

This behavior is ballistic, meaning that its actions continue even after its activation condition is no longer satisfied. This is useful for behaviors which might otherwise exhibit oscillatory activity. The idea is to take the system well out of the current conditions, instead of placing it at the boundary where it can easily re-enter the previous, undesirable state. In this case, the robot will keep moving backwards even after the obstacle is out of range. This way it is more likely to reach a safe position where a higher-level behavior can take over.

B. Collision avoidance (CA)

Collision avoidance attempts to prevent collisions at an earlier stage by stopping the robot and turning away from the nearest reflector until no further reflectors are in range. This ensures that after its execution, the robot has an unhindered path in front of it:

$$\begin{aligned} C &= (n < T_d \wedge (s_L^{env}[n] > T_A \vee s_R^{env}[n] > T_A)) \\ V_l &= 0 \\ V_a &= V_a^{max} \cdot \text{sgn}(\Delta r) \end{aligned}$$

with T_d corresponding to 0.2 m and V_a^{max} equal to 1 rad/s.

This behavior is also ballistic; the robot will keep turning in the same direction until no more reflectors are within an extended range. This is to prevent oscillations which might otherwise occur, for example in corners. Ballistic behaviors can still be interrupted by lower-level behaviors. In this case the CD behavior can interrupt the CA behavior when an object comes too close while the robot is turning.

C. Obstacle avoidance (OA)

Obstacle avoidance enables the robot to follow smooth trajectories around objects. It will adjust its translational speed proportional to the distance to the closest reflector. Its rotational speed calculated inversely proportional to the distance to the closest reflector, as well as to the difference between the left and right reflectors. This allows the system to curve its path around distant obstacles and to pass through doors:

$$\begin{aligned} C &= (n < T_d \wedge (s_L^{env}[n] > T_A \vee s_R^{env}[n] > T_A)) \\ V_l &= V_l^{min} + (V_l^{max} - V_l^{min}) \cdot \sqrt[3]{\frac{d_r}{T_d}} \\ V_a &= V_a^{min} \cdot \text{sgn}(\Delta r) + (V_a^{max} - V_a^{min}) \cdot \left(\frac{\Delta r}{T_d} \cdot \frac{T_d - d_r}{T_d}\right) \end{aligned}$$

with T_d corresponding to 1 m, V_l^{min} equal to 0.05 m/s, V_l^{max} equal to 0.2 m/s, V_a^{min} equal to 0.05 rad/s and V_a^{max} equal to 0.8 rad/s.

One thing to note is that the system cannot clearly distinguish between a single reflector straight ahead and two separate reflectors at exact opposite angles of the normal. Although theoretically this would cause the system to drive directly at objects when encountered straight ahead, the metastability of this situation is easily disturbed and the robot will divert to one side of the obstacle. Similarly, doors and corners appear identical to the system, both as two equally-spaced reflectors. Because the behavior is made to pass through doors, it will




Fig. 5: Overview of the system used for autonomous navigation. (a) The physical system consisting of the Pioneer mobile robot (bottom, red), Hokuyo laser rangefinder (middle, blue) and custom radar sensor (top, white). The Intel NUC PC is located behind the laser rangefinder. (b) Schematic representation of the hardware and software components of the system. It should be noted that lidar is used only for gathering ground truth data and not for autonomous navigation or SLAM.

attempt to perform the same action at a corner. Although this would cause it to get stuck in the very tip of the corner, the system relies on the lower-level behaviors to take over and put the robot back on a suitable track.

D. Corridor following (CF)

Corridor following tries to steer the robot through its environment on stable paths. It does this by balancing the amount of change in energy of both channels. This is achieved converting a signal's waveform to its spectrogram representation. Next, the difference between the current spectrogram and that of the previous time step is calculated for each channel. Finally, the angular velocity is based on the difference in change between the left and right channel.

$$\begin{aligned}
 C &= (|\Delta s| > T_A) \\
 V_l &= V_l^{max} \\
 V_a &= V_a^{max} \cdot \sqrt{|\Delta s| - T_A} \cdot \text{sgn}(\Delta s)
 \end{aligned}$$

with T_d corresponding to 4 m, V_l^{max} equal to 0.2 m/s and V_a^{max} equal to 0.5 rad/s.

When the robot is driving in an environment such as a corridor, it might end up on a path close to one of the walls. This will cause echoes to vary strongly on the side facing the wall, because any irregularities of the surface pass by quickly. On the other facing side these echoes will vary less because of the distance between the surface and the sensor. As a result, the robot will turn away until the change in the spectral energy is equal and the path of the robot is centered. A side effect is that when one of the walls is rougher than the other, i.e. causes more change in spectral energy, the stable path of the robot will not coincide with the centerline of the corridor. Instead it will lie parallel to it, but farther away from the rough surface. This behavior is based on the concept of optical flow, which is present in both insect and mammal vision [25]. In essence, optical flow is the extraction of information about movement from a series of consecutive visual images. It can be used to estimate the egomotion of a system, and to regulate speed or correct for drift in closed loop control. This principle also

applies to active sensing mechanisms such as sonar and radar, where it is called acoustic, or echoic, flow [26]. By analyzing the continuous variation of echo signals through time, similar functionality to optical flow can be achieved.

E. Straight drive (SD)

Straight drive is the default behavior when none of the other behaviors are activated. It moves the robot forward in its current direction at a constant speed until another behavior takes over control.

$$\begin{aligned}
 V_l &= V_l^{max} \\
 V_a &= 0
 \end{aligned}$$

The path of the robot might not be completely straight, as any misalignment of the wheels causes drift over time.

IV. EXPERIMENTAL RESULTS

In this section we demonstrate the performance of our system in two distinct environments. For this purpose, we use the robotic setup shown in 5. This setup is composed of a Pioneer P3DX mobile platform, Intel NUC NUC5i7RYB PC, Hokuyo UBG-04LX-F01 laser rangefinder, and the modified Flat Earth Ancho radar kit. The NUC hosts a MATLAB instance which interfaces with the radar sensor and executes the control architecture, and also hosts several ROS nodes which interface to the mobile robot and generate ground truth data. The lidar sensor serves the sole purpose of generating ground truth data with which our system can be validated. This ground truth consists of a metric map of the environment, generated with the GMapping ROS-module, and the exact trajectory driven by the robot, produced by the AMCL ROS-module [27], [28]. These are not used in autonomous navigation or SLAM, but only for verification and visualization purposes.

In both cases the robot was placed in part of a real-world office environment where some obstacles were added to increase the difficulty of navigation. The first environment is a corridor with added obstacles as shown in Figure 6. This layout serves as a suitable testbed to determine how the system responds when forced to pass close to objects, how quickly its path stabilizes afterwards, and how it handles dead ends. The second environment is a hallway with benches, vending machines and open spaces as shown in Figure 7. This environment has more variability in both space as well as the type and size of objects present. It allows us to observe how the system behaves and stabilizes when less constrained in its movement.

A. Autonomous navigation

A comparison of the performance of the different stacks of behaviors within a corridor environment is shown in Figure 6, for which we will discuss each stack separately. The produced trajectories each have a length of approximately 186 m, which corresponds to three full passes through the environment. Panel (a): the movement of the robot operated by the lower-level behaviors (CD, CA) is very basic. The robot only changes its trajectory when it encounters an object right in front of




Fig. 6: Trajectories of the different stacks of behaviors in an office corridor with added obstacles. (a) The path of the robot with only the lowest-level behaviors active, namely SD, CD and CA. (b) The resulting path when adding OA to the previous behaviors. (c) The path taken by the full-stack architecture, i.e. SD, CD, CA, OA and CF.




Fig. 7: Trajectories of the different stacks of behaviors in an office hallway with benches, vending machines and open space. (a) The path of the combined SD, CD and CA behaviors. (b) The path when including the OA behavior. (c) The path of the full-stack architecture with CF as well as the previous behaviors.

it, making for very erratic movements. Panel (b): by adding the OA behavior to the stack, the movements become more fluent. This can be attributed to the system's capability of responding to obstacles from farther away and to the controller now including a non-linear component. We can observe that the robot curves around almost all obstacles, except when approaching the rightmost one. This is likely because it is centered in the corridor (see explanation in Section III-C). Furthermore, near obstacles the robot tends to follow the same path for each pass in a specific direction. Only in the open spaces is there more variability in its path. Panel (c): when including the CF behavior the trajectory becomes even more smooth. This is an effect of the behavior being more sensitive to smaller perturbations; it takes into account all changes in spectral energy instead of looking only at the main reflectors. This also results in the forward path through the environment now coinciding with the return path.

Next we take a look at the performance of the system in the second environment (Figure 7). These trajectories have a length of approximately 225 m, resulting in about four




Fig. 8: Distribution of the angular velocities for each pose of a trajectory. Subfigures (a) to (c) show histograms of the angular velocities for the trajectories of Figures 7a to 7c respectively.

and a half passes through the environment. In the previous environment, the stability in the trajectory of the full-stack system could be attributed to the highly constrained space. In this environment the robot has more leeway and thus we can determine if the stability is inherent to the system itself or not. Furthermore, we can establish how well the system handles a different environment which contains objects of varying size and materials. Panel (a): using the lower-level set of behaviors, the system behaves in a very similar way to the previous environment; arbitrary movement throughout the available space, no overall pattern can be discerned. Panel (b): by adding the OA behavior, the trajectory becomes more regular. However, there is still variability in the sections where the path is less restricted. Furthermore, the system has to recover from fewer near-collisions with the environment, because it is able to start maneuvering at a farther distance. Panel (c): when including the CF behavior in the stack, the overall trajectory of the system stabilizes. Consecutive passes of the same sections overlap highly for the majority of the trajectory. This overlap can again be credited to the fact that the CF behavior takes into account the information of the full surroundings within line of sight. After encountering an obstacle, the position of the robot will quickly converge to the most stable path in the environment, without depending much on the angle at which the obstacle was encountered.

To quantify the smoothness of a trajectory, we have chosen the distribution of the angular velocities for each pose of that trajectory as metric. By looking at Figures 6 and 7, we can intuitively estimate the smoothness of each trajectory. Trajectories with only the low-level behaviors enabled, consist of long straight parts with sharp turns in between. The trajectories of the full-stack behavior on the other hand exhibit more curvature and less discontinuities. The distribution of angular velocities is visualized in Figures 8a to 8c, for the trajectories of Figures 7a to 7c respectively. This shows that the first trajectory has more abrupt turns according to the higher spread of the distribution. For the third trajectory the distribution is more concentrated, indicating more fluent turns overall. This metric can be further condensed by taking the standard deviation of the distributions, which are 0.255, 0.183, and 0.137 respectively.

In Figure 9 the active behavior for each position of the given




Fig. 9: Visualization of the active behaviors for each position of the trajectory in Figure 7c. Behaviors are color-coded and higher-priority behaviors are plotted on top in case of overlap.

trajectory is shown. In this case the full stack of behaviors was enabled on the system; CD was active for 0.4% of the trajectory, CA for 0.8%, OA for 12.7%, CF for 34.2% and SD for 51.9%. The reason for SD being the most active behavior is that the environment has several parts of open space; here there are no obstacles and the spectral variation is low. CF is the most active dynamic behavior. It occurs particularly in curved parts of the trajectory, because of the difference in spectral change on each side of the sensor. The system will activate to make the robot turn such that this change is balanced, thus following the curve. When the angular correction of the CF behavior is not enough to steer the robot away from an obstacle or when an obstacle suddenly appears in line of sight, the OA behavior will take over. Because it causes the robot to slow down considerably when close to an obstacle, the system has enough time to turn and avoid the obstacle. The sparse occurrence of the lower-level behaviors can be explained by the fact that high-level behaviors prevent the system from entering critical situations where they are necessary. The CA behavior only takes over in the upper right corner of the map, once for each loop of the trajectory. This can be attributed to the robot first passing through an open space where no correction are made, until the left wall appears in sight. By this time, the system can't respond fast enough to smoothly turn the robot around. The CD behavior is activated only once for the entire trajectory, where the robot managed to squeeze itself between two objects in the upper left corner of the map.

The presence of the stable paths we observe with the addition of the CF behavior can be regarded as an emergent property of the system. None of the behaviors are explicitly programmed to adhere to the same path when traversing a section multiple times. However, we can expect traces through the environment for which the difference in change of spectral energy (Δs) is minimal. As long as the the lower behaviors ensure the robot passes near these minima, the dynamics of the system will cause it to end up on the same path, even in the more open areas.

B. SLAM

During the experiments on autonomous navigation, we recorded the raw radar signals and odometry of the robot. Using this data we can perform simultaneous localization




Fig. 10: Topological SLAM using the sensor data and odometry gathered during the trajectory of Figure 7c. (a) The trajectory generated by path integration of the odometry of the robot. Cropped to the map boundaries. (b) The ground truth map and trajectory. (c) The trajectory generated by the RadarSLAM algorithm [11].




Fig. 11: Histogram of the absolute and cumulative localization errors compared to the ground truth of the trajectories in Figure 10. (a) Localization error of the trajectory generated by odometry as seen in Figure 10a. (b) Localization error of the trajectory generated by RadarSLAM as seen in Figure 10c.

and mapping with RadarSLAM, a biologically-inspired SLAM solution for radar developed in our previous work [11]. In summary, this system works by building a database of the spectrograms observed at each visited position. By comparing the spectrogram at the current position to those in the database, while taking into account the recent odometry of the robot, the system is able to determine the most probable location of the robot. If necessary the system corrects any error between the previous estimate and the updated one, which is called a loop closure. In this case the system also applies graph relaxation, which adjust previous poses to better match the path leading up to the corrected pose. As RadarSLAM is a graph-based SLAM algorithm, the system outputs a topological map. This means that the map respects adjacency of locations, but does not preserve absolute distance or direction. This is unlike a metric map, which can be regarded as a scaled representation of the real environment. A more detailed description on RadarSLAM can be found in the original paper [11].

For the implementation of RadarSLAM used in this paper, the spectrograms of the signals at the left and right receiver are concatenated to form the local view template V :

$$V = (\mathcal{S}_L | \mathcal{S}_R)$$

with V having size $M \times L$ and L being equal to $2K$. The distance metric d_V is taken as the inverse of the correlation

coefficient ρ between two template vectors V and W :

$$\rho(V, W) = \frac{1}{N-1} \sum_{n=1}^N \left(\frac{V_n - \mu_V}{\sigma_V} \right) \left(\frac{W_n - \mu_W}{\sigma_W} \right)$$

$$d_V = 1 - \rho(V, W)$$

with μ the mean of a template, σ its standard deviation, and N equal to $M \cdot L$. The use of this type of metric makes RadarSLAM a direct SLAM method.

One of the main advantage of the current setup compared to that of the previous RadarSLAM paper, is the improved design of the sensor. The use of two receivers at different angles allows the system to capture more information about the environment, which makes place recognition less ambiguous. When applied to the trajectory of Figure 7c the system produces accurate localization results, as can be seen in Figure 10. The odometry exhibits a large drift over time which makes the original trajectory unrecognizable. However, RadarSLAM is able to successfully approximate the real trajectory of the robot. The same conclusion can be drawn from the metrics presented in Figure 11. For the odometry, the average error is 6.24 m, while the maximum error is 18.54 m. The SLAM algorithm on the other hand has an average error of 0.35 m and a maximum error of 1.4 m. This makes the SLAM results for this trajectory on average 17 times more accurate than those of the raw odometry. The reason for this improvement is twofold; first, the odometry exhibits a high angular bias which causes it to accumulate a large error over the course of the trajectory. This bias is negated by the continuous corrections done by the SLAM algorithm. Second, because the system traverses stable paths, it visits roughly the same locations on each pass through the environment. This means that a large portion of the spectrograms gathered in the first pass can be matched to spectrograms encountered in consecutive passes, which greatly facilitates localization.

V. CONCLUSION

This paper describes a system for autonomous navigation using biomimetic radar sensor. We show that both the hardware and signal processing aspects of the system exhibit definite similarities to the mechanisms used by an echolocating bat, in this case the common big-eared bat (*Micronycteris microtis*). The use of a single emitter in combination with two receivers at an opposite angle of each other gives rise to the interaural intensity and time differences, which are also present in biological echolocation. These features can be used to locate reflectors in the horizontal plane, thus allowing the system to purposefully avoid obstacles. The signal generated by the radar sensor and the call produced by the *M. microtis* also share several key properties. Although the former is a Gaussian-modulated sine wave and the latter a frequency-modulated sweep, both signals are very short pulses which exhibit a wide frequency band. This type of signal has the advantage that it has a high spatial accuracy and a large spectral content. The echo-related transfer function of the radar, which is determined by the combined radiation patterns of the emitter and receivers, also bears resemblance to its biological counterpart; a main lobe with high directivity, which varies with the frequency of

the signal. However, the beam of the antenna has a higher drop-off in intensity and less overall variation, which reduces the ability of the system to use the spectral cues to their full extent. Having gained this knowledge, we intend to improve the variability of the ERTFs by optimizing the antennas in future work. This will allow the system to make better use of spectral cues generated by the reflectors in the environment, which is required to determine the direction of obstacles in the vertical plane.

We have also demonstrated the performance of the autonomous navigation capabilities in two distinct real-world environments. We can observe significant improvements between the different stacks of behavior, as expected; the combination of the straight drive, collision detection and collision avoidance behaviors results in arbitrary bouncing around the environment, without much of a pattern to be discerned. When we add the obstacle avoidance behavior to this stack, the movement of the robot becomes much smoother and an overall pattern starts to appear in the trajectory. Finally, by including the corridor following behavior, movement is even further smoothed and the system starts to exhibit stable paths throughout most sections of the environment.

Furthermore, the system is capable of successfully performing SLAM using the RadarSLAM algorithm developed in previous work. Even though the raw odometry exhibits a high bias in its angular accuracy, the system is still able to localize itself with an average error of 0.35 m.

REFERENCES


- [1] K. Bengler, K. Dietmayer, B. Farber, M. Maurer, C. Stiller, and H. Winner, "Three decades of driver assistance systems: Review and future perspectives," *IEEE Intelligent Transportation Systems Magazine*, vol. 6, no. 4, pp. 6–22, 2014.
- [2] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. N. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer, M. Gittleman, S. Harbaugh, M. Hebert, T. M. Howard, S. Kolski, A. Kelly, M. Likhachev, M. McNaughton, N. Miller, K. Peterson, B. Pilnick, R. Rajkumar, P. Rybski, B. Salesky, Y.-W. Seo, S. Singh, J. Snider, A. Stentz, W. Whittaker, Z. Wolkowicki, J. Ziglar, H. Bae, T. Brown, D. Demitrish, B. Litkouhi, J. Nickolaou, V. Sadekar, W. Zhang, J. Struble, M. Taylor, M. Darms, and D. Ferguson, "Autonomous driving in urban environments: Boss and the urban challenge," *Journal of Field Robotics*, vol. 25, no. 8, pp. 425–466, 2008.
- [3] H. U. Schnitzler, C. F. Moss, and A. Denzinger, "From spatial orientation to food acquisition in echolocating bats," *Trends in Ecology and Evolution*, vol. 18, no. 8, pp. 386–394, 2003.
- [4] D. Von Helversen and O. Von Helversen, "Acoustic guide in bat-pollinated flower," *Nature*, vol. 398, no. 6730, pp. 759–760, 1999.
- [5] J. Steckel and H. Peremans, "Acoustic Flow-Based Control of a Mobile Platform Using a 3D Sonar Sensor," *IEEE Sensors Journal*, vol. 17, no. 10, pp. 3131–3141, 2017.
- [6] —, "BatSLAM: Simultaneous Localization and Mapping Using Biomimetic Sonar," *PLoS ONE*, vol. 8, no. 1, pp. 1–11, 2013.
- [7] J. W. Starr and B. Y. Lattimer, "Evaluation of Navigation Sensors in Fire Smoke Environments," *Fire Technology*, vol. 50, no. 6, pp. 1459–1481, 2014.
- [8] C. C. Chen, "Attenuation of Electromagnetic Radiation by Haze, Fog, Clouds, and Rain," RAND Corporation, Santa Monica, CA, Tech. Rep., 1975. [Online]. Available: <http://www.rand.org/pubs/reports/R1694.html>
- [9] R. A. Brooks, "A robust layered control system for a mobile robot," *IEEE Journal on Robotics and Automation*, vol. 2, no. 1, pp. 14–23, 1986.
- [10] R. C. Arkin, *Behavior-based robotics*. MIT press, 1998.
- [11] G. Schouten and J. Steckel, "RadarSLAM: Biomimetic SLAM using ultra-wideband pulse-echo radar," in *2017 International Conference on Indoor Positioning and Indoor Navigation, IPIN 2017*, vol. 2017-January, no. September, 2017, pp. 1–8.

- [12] *Ancho Data Sheet*, Flat Earth Inc., 2015. [Online]. Available: <https://www.xethru.com/community/resources/ancho-data-sheet.49/>
- [13] *X2 Datasheet*, Novelda AS, 2015. [Online]. Available: <https://www.xethru.com/community/resources/x2-datasheet.24/>
- [14] *BeagleBone Black System Reference Manual*, Texas Instruments, 2013. [Online]. Available: https://cdn-shop.adafruit.com/datasheets/BBB_SRM.pdf
- [15] *Novelda 6.0 - 8.5 GHz Sinuous Antenna*, Novelda AS, 2011. [Online]. Available: <https://www.xethru.com/community/resources/documentation-suite.43/>
- [16] *Getting Started with Salsa Cilantro*, Flat Earth Inc., 2017. [Online]. Available: <https://store.flatearthinc.com/products/cilantro-switching-capacitor>
- [17] B. Grothe, M. Pecka, and D. McAlpine, "Mechanisms of Sound Localization in Mammals," *Physiological Reviews*, vol. 90, no. 3, pp. 983–1012, 2010.
- [18] R. A. Holland, "Echolocation signal structure in the Megachiropteran bat *Rousettus aegyptiacus* Geoffroy 1810," *Journal of Experimental Biology*, vol. 207, no. 25, pp. 4361–4369, 2004.
- [19] J. Steckel and H. Peremans, "A novel biomimetic sonarhead using beamforming technology to mimic bat echolocation," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 59, no. 7, pp. 1369–1377, 2012.
- [20] D. Vanderelst, F. De Mey, H. Peremans, I. Geipel, E. Kalko, and U. Firzlaff, "What Noseleaves Do for FM Bats Depends on Their Degree of Sensorial Specialization," *PLoS ONE*, vol. 5, no. 8, p. e11893, aug 2010.
- [21] J. Marple L., "Computing the discrete-time "analytic" signal via FFT," *Signal Processing, IEEE Transactions on*, vol. 47, no. 9, pp. 2600–2603, 1999.
- [22] J. Allen, "Short term spectral analysis, synthesis, and modification by discrete Fourier transform," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 25, no. 3, pp. 235–238, jun 1977.
- [23] L. Robles and M. a. Ruggero, "Mechanics of the Mammalian Cochlea," *Physiological Reviews*, vol. 81, no. 3, pp. 1305–1352, jul 2001.
- [24] J. Steckel and H. Peremans, "Acoustic Flow-Based Control of a Mobile Platform Using a 3D Sonar Sensor," *IEEE Sensors Journal*, vol. 17, no. 10, pp. 3131–3141, 2017.
- [25] M. V. Srinivasan, "Visual control of navigation in insects and its relevance for robotics," *Current Opinion in Neurobiology*, vol. 21, no. 4, pp. 535–543, aug 2011.
- [26] R. Müller and H.-U. Schnitzler, "Acoustic flow perception in cf-bats: Properties of the available cues," *The Journal of the Acoustical Society of America*, vol. 105, no. 5, pp. 2958–2966, may 1999.
- [27] G. Grisetti, C. Stachniss, and W. Burgard, "Improved techniques for grid mapping with Rao-Blackwellized particle filters," *IEEE Transactions on Robotics*, vol. 23, no. 1, pp. 34–46, feb 2007.
- [28] C. Kwok, D. Fox, and M. Meila, "Adaptive real-time particle filters for robot localization," *2003 IEEE International Conference on Robotics and Automation (Cat. No.03CH37422)*, vol. 2, pp. 2836–2841, 2003.



he develops novel radar architectures inspired on the echolocation system of bats.

Jan Steckel graduated in electronic engineering at the University College Karel de Grote in Hoboken in 2007 and in 2012 his doctoral degree on array signal processing. In 2015 he became a tenure track professor at the University of Antwerp in the Constrained Systems Lab where he researches sensors, sensor arrays and signal processing algorithms using an embedded, constrained systems approach. His current research interests are focused on industrial applications of ultrasonic sensing and on the development air-coupled ultrasonic imaging sensors. Furthermore,



Girmi Schouten obtained his bachelor's degree in Applied Information Technology in 2014 at the Artesis Hogeschool, Antwerp. He went on to complete the Master of Electronics and ICT Engineering Technology at the University of Antwerp in 2016. Currently, he is working towards his doctoral degree at the University of Antwerp under Prof. Dr. Ing Jan Steckel on a grant of the FWO (Research Foundation Flanders). His present research focuses on transferring the research group's knowledge of sonar to the domain of radar. This includes ultra-

wideband radar signal processing and antenna design, as well as developing control architectures and SLAM algorithms.