Low-power multichannel spectro-temporal feature extraction circuit for audio pattern wake-up

Dinko Oletic, Vedran Bilas
Faculty of Electrical Engineering and Computing
University of Zagreb
Unska 3, 10000 Zagreb, Croatia
vedran.bilas@fer.hr

Michele Magno, Norbert Felber, Luca Benini
Dept. of Information Technology and Electrical Eng.
ETH Zürich
Gloriastrasse 35, 8092 Zürich, Switzerland
lbenini@iis.ee.ethz.ch

Abstract*—In many distributed sensing applications, continuous sensor monitoring requires processing with a significant energy footprint, which hinders autonomous operation and battery lifetime of sensor nodes. In our research we explore the power savings gained by splitting the hardware architecture for continuous monitoring into two stages: an always-on ultra-low-power mixed-signal wake-up circuit placed near the sensor, performing coarse recognition (e.g. wake-up circuit) and waking up the main digital processing unit only on event detection. This enables for activation of energy-hungry digital processing only at the rate of event occurrence without penalising responsiveness and monitoring continuity. We focus on the wake-up circuit performing recognition of spectro-temporal audio patterns, consisting of spectro-temporal feature extraction, and the classification sub-circuits. We propose a novel design of the feature extraction circuit. It consists of a spectral decomposition multi-channel analog band-pass filter bank, implemented in generalized impedance converter topology (GIC), and the bank of passive channel detectors for measuring the intervals of in-band signals. Experimental filter characterization demonstrated the benefits of proposed filtering topology for low-power applications in the audio frequency range even with operational amplifiers of very limited bandwidth. Detector's response was verified in multi-channel environment. Preliminary analysis showed power consumption ranging from 10.5 to 13.5 µW per channel using off-the-shelf components.

Keywords—event monitoring, energy efficiency, audio pattern detection, wake-up circuit, generalized impedance converter.

I. Introduction

In many distributed sensing applications, sensor nodes (consisting of sensors and digital signal processors) are required to operate autonomously for prolonged time periods ranging from days to years, often in the environment with a limited power supply at their disposal [1]. Capacity of battery powered devices is typically limited by size and cost, typically to milliwatt-hours, defining the ceiling of their total average power to microwatts, in order to reach targeted lifetime. Present-day sensors, acquisition circuits, and DSPs often exceed this power-budget in many application domains [2]-[4], making event monitoring challenging to implement.

A common approach of increasing the sensor node's lifetime, is employment of dynamic power management mechanisms on its scalable hardware [5]. Such mechanisms typically entail fixed or adaptive duty-cycling schemes, switching between active and low-power inactive states, to minimize the activity interval of power hungry hardware [6]. Event detection implemented by digital signal processing traditionally implies periodic signal sampling, storing and processing, at some sampling frequency sufficient for the particular application, to avoid missing the event, as shown in the upper portion of Fig. 1. However, the rate of occurrence of the event of interest is usually very low (rare). Thus, even by employing aggressive power management strategies, such as duty-cycling to lowest-power states between successive samples, or shortening processing time by optimizing pattern recognition algorithms [7], much energy is wasted on sampling, storing and processing irrelevant information.

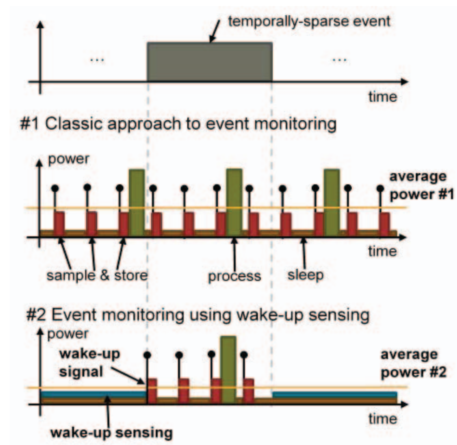

Fig. 1. Comparison of event monitoring/detection concepts. Up: continuous signal sampling and processing. Down: two-stage wake-up architecture.

In order to mitigate the problem of the total power spent for event detection, we propose an alternative, two-stage architecture consisting of: (1) wake-up circuit with feature extraction, and (2) main digital processing unit (microcontroller or digital signal processor). The wake-up circuit is an ultra-low-power, but always-on block, continuously monitoring the sensor signal. The circuit operates either in the analog or mixed signal domain, and provides a coarse recognition of some pattern related to occurrence of the monitored event. Upon recognition, it outputs a digital wake-up signal engaging the second stage, implementing the classical signal sampling and digital signal processing for a more thorough pattern recognition. This enables the energy-

978-3-9815370-6-2/DATE16/

hungry main processor to be completely turned-off until reception of the wake-up signal. In case of low-enough consumption of the wake-up circuit, with respect to the event occurrence rate, wake-up architecture may bring substantial power savings (Fig. 1, lower portion).

We focus on a circuit performing wake-up upon recognition of predefined audio patterns, for event monitoring in applications related to urban environments, such as traffic, assisted-living, health and well-being (horns, sirens, ringing, ticking, engine sounds, whistles, vocalization, etc.). There, audio signals of interest resemble a known spectro-temporal structure, and thus may be classified based on their spectro-temporal features. In this paper we show a prototype implementation of analog circuit design performing multi-channel extraction of spectro-temporal features: intervals of signal containing spectral content in the preset frequency bands. The main contributions of the paper are:

- design of the programmable filtering stage using generalized impedance converter topology enabling for usage of micro-power operation amplifiers,
- baseline independent passive channel detector design,
- analysis of the state-of-the-art boundaries of power-consumption achievable with discrete off-the-shelf components.

The rest of the paper is organized as follows: section II describes the related work; section III overviews architecture of our low-power audio pattern wake-up circuit, section IV focuses on the design of the feature extraction sub-circuit. Section V demonstrates experimental results, and section VI concludes the paper.

II. RELATED WORK

Reducing energy consumption of sensor nodes (in particular for long term monitoring) has been a topic of intensive research in recent years. The most common energy-saving method for these devices is duty-cycling [5]. These methods can be roughly divided into two categories: synchronous and asynchronous. The major limitations of synchronous methods is time synchronization which can lose important events, and the trade-off between energy reduction and sleep periods. Asynchronous methods, on the other hand, allow sensor nodes to operate independently, by switching between active and sleep mode, at the rate of event occurrence [11]-[10].

In long-term continuous audio monitoring using low-power microphones, asynchronous methods are preferred due to burden of signal processing. They pursue the improvement in energy efficiency by using an analog wake-up stages for preliminary event detection. Simple examples of acoustic wake-up circuits, featuring wake-up on magnitude-based detection of acoustic activity, are shown in [11] and [12], where effectiveness of the approach is supported by extensive simulation of power-savings resulting from inclusion of the wake-up circuit. More advanced design, from the field of human speech detection [13], demonstrated the integrated wake-up circuit within a power-budget of 500 μW.

Ultra-low-power analog designs for audio signal detection, integrated on chip, are typically reported in wearable designs of medical devices, such as hearing impairment aids (cochlear implants etc.). However, design ideas proposed in this field are universally applicable. They show that in silicon, on-chip spectral decomposition by analog band-pass filtering may be implemented for as low as 4.5 μW per channel [14]. A 63.6 μW, high-precision, programmable, 16-channel analog filtering bank is shown in [15]. Adaptive single channel audio magnitude detector was implemented in power budget of 1.1 μW [16]. Finally, a complete, analog bionic ear (cochlea) chip with total consumption of 211 μW is shown in [17].

Mentioned designs entail custom silicon chip design. In many applications, design with discrete components may be favored in order to shorten the development cycle, facilitate scalability, and reduce development costs. However, power consumption boundaries of the audio pattern recognition wake-up circuit, featuring state-of-the-art discrete off-the-shelf components, are not clear. Thus, we address them in this paper, for a circuit featuring functionalities comparable to those shown in [14]-[16].

III. LOW-POWER AUDIO PATTERN RECOGNITION WAKE-UP ARCHITECTURE

In this section we give a high-level architectural overview of the ultra-low-power always-on circuit, performing low-power recognition of an acoustic pattern in the analog/mixed signal domain. Upon recognition, the circuit generates a digital signal which wakes-up the main digital signal processing stage. Generally, wake-up circuit can be broken down to two parts: (1) feature extraction circuit, and (2) classification circuit. Circuit operation is illustrated in an example shown in the upper portion of Fig. 2, depicting a spectro-temporal (time-frequency) decomposition of an audio event, obtained by the short-term Fourier transform (STFT spectrogram).

The spectrogram shows the change in the “ticking” sound at the road crossing notifying pedestrians about the traffic light transition from red to green. It can be seen that the event may be uniquely characterized, by observing the spectrogram's magnitudes, contained within the ordered sequence of regions highlighted in rectangles (e.g. magnitude within frequency band of width B_1, centered at frequency f_{c1}, of duration T_1, followed by the pause T_2 etc.). Therefore, the feature extraction circuit, shown in lower part of Fig. 2, extracts a two-dimensional, spectro-temporal feature-set.

Following the spectral decomposition, a band-pass filtered signal is fed to a bank of channel detectors for extraction of temporal features. Depending on the application, the circuit may implement extraction of the different quantities: time-instants of occurrence, duration, magnitude (envelope, amplitude), accumulated energy etc. [17] of the signal exceeding some absolute or relative threshold. Passive implementations of the channel detectors are preferred for lower power consumption. The thresholds constants (e.g. C_1, C_2, C_3 etc.) should be dynamically programmed by the classification circuit.

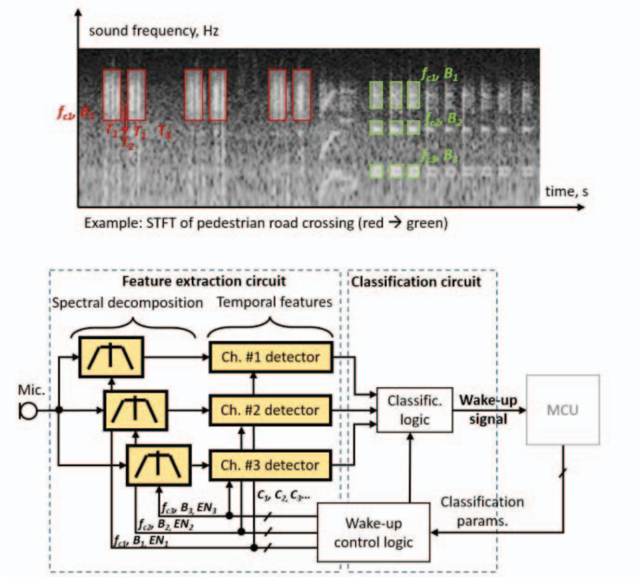

Fig. 2. Up: Spectro-temporal content of the example audio event. Regions of interest enabling the recognition of the event are highlighted red/green. Down: High-level overview of the always-on circuit performing MCU/DSP wake-up based on recognition of the spectro-temporal pattern in audio signal from microphone.

Binary classification is performed upon the extracted features by the classification logics block. Positive outcome results in activation of the digital wake-up signal, triggering the MCU from sleep state. One major requirement for classification block is low-power mixed-signal or digital implementation. Given the sequential nature of the audio patterns, in the most simplistic case, classification might be implemented by the sequential template matching, similar to [18]. However, more complex classification methods may be applied, such as hidden Markov model [19] etc.

Whole pattern recognition process is governed by a common wake-up control logic, which dynamically sets the parameters for the feature extraction, and classification blocks. Parameters are preloaded from MCU and stored in the wake-up control logic block.

As the wake-up circuit is supposed to be always on, aside from the pure event detection performance, power consumption is its main design requirement. In the rest of the paper, we focus on the efficient design of the spectro-temporal feature extraction circuit (highlighted yellow in Fig. 2).

IV. Design of the Spectro-Temporal Feature Extraction Circuit

In this section we present the details of the proposed architecture for an ultra-low-power spectro-temporal feature extraction circuit from Fig. 2, consisting of a bank of channel filters and detectors, using discrete off-the-shelf components. Although multiple channels are required in parallel for simultaneous spectro-temporal decomposition, the principle of operation is explained for a single channel. The schematic is shown in Fig. 3 and it consists of a channel filter, and a channel detector.

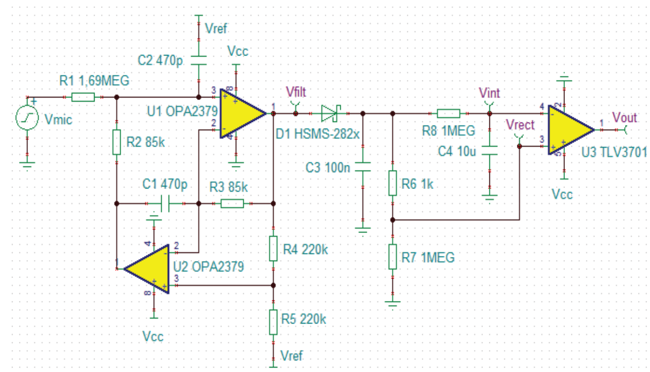

Fig. 3. Schematics of a single channel of the feature extraction circuit performing spectral decomposition (band-pass filtering) and measurement of duration of in-band spectral content.

A. Channel Filter Design

In filter design with discrete operational amplifiers, power-budget is dominated by (1) choice of amplifier model (its supply current) and (2) filtering topology.

Amplifier's speed, defined by its gain-bandwidth product and slew-rate parameters, limits the filter's operating frequency range [20]. This manifests through the discrepancies between theoretical and real filter response (errors of gain, central frequency, and Q-factor), which increase with frequency. Moreover, non-idealities worsen with amplifier types/models featuring lowest supply current, limiting their usability typically to sub-kHz applications.

However, choice of filtering topology can minimize effects of the operational amplifier's non-idealities. Therefore, generalized impedance converter (GIC) filtering circuit topology [20] was chosen for implementation of channel filters. The band-pass GIC filter is shown in Fig. 3 between voltage nodes V_{mic} and V_{filt}, comprising of two operational amplifiers, U_1, U_2, two capacitors C_1, C_2, and five resistors $R_1 \ldots R_5$. In comparison to classical Sallen-Key or multiple-feedback topology, GIC topology enables for:

- inherent compensation of the operational amplifier's non-idealities by setting: (1) the equal time-constants $C_2 \cdot R_2 = C_1 \cdot R_3$, (2) choosing the equal resistor values $R_4 = R_5$ [20], and (3) using dual-package amplifiers featuring similar characteristics. This lowers the power consumption, by enabling the usage of operational amplifiers with lowest supply current (in 1 µA range).
- independent setting of filtering parameters: central frequency f_c set by two resistors equal in value, and pass-bandwidth B, using the single resistors,
- predetermined, constant voltage gain of 2 (+6 dB), at central frequency f_c, independent of Q-factor and other parameters.

Filter's transfer characteristic, written with resistors expressed by their respective conductance $G_1 \ldots G_5$ (for brevity), is given by (1):

$$H(s)=\frac{V_{filt}}{V_{mic}}=\frac{G_1(sC_1G_4+sC_1G_5)}{s^2C_1C_2G_4+sC_1G_4G_1+G_5G_3G_2}. \quad (1)$$

Filtering parameters are programmed by resistors R_1, R_2 and R_3. Remaining resistors R_4, R_5 can be chosen freely so far they are kept equal in value. Capacitors $C_1 = C_2 = C$ are also fixed. Central filtering frequency f_c is set by choosing $R = R_2 = R_3$ as per (2):

$$R=R_2=R_3=\frac{1}{2\pi f_c C}. \quad (2)$$

Pass-band width B is set by R_1, according to (3):

$$R_1=\frac{f_c}{B}R. \quad (3)$$

Filtering amplifiers U_1 and U_2 are implemented using micro-power operational amplifiers OPA2379 (Texas Instruments) featuring supply current of 2.9 µA per channel at 1.8 V, gain-bandwidth product of 90 kHz, slew rate of 30 mV/µs and common-mode rejection ratio of 100 dB [21].

To support operation with the DC-biased microphone's output signal at around 450-550 mV, the filter is implemented in single-supply version, with supply voltage V_{cc} set to 1.8 V, and reference (common mode) voltage V_{ref} equaling $V_{cc}/2$. The value of R_1 is compensated for the microphone's large output impedance. For the MEMS microphone model used ICS-40310 (InvenSense), it typically ranges from 4.5 to 5.5 kΩ [22].

B. Channel detector

Temporal features are extracted by the channel detector, connected at the output of the filter, V_{filt}. It tracks the intervals of duration of the in-band signal, in which absolute value of the signal's magnitude is larger than the predefined threshold. In this case it keeps V_{out} = 'high', otherwise V_{out} ='low'.

Implementation is analogous to the address decoder for a wake-up radio [23], with the alternative design of the baseline tracking, used to avoid dependence on fixed thresholds. The detector is implemented as a passive design to minimize the overall power consumption, by a cascade of the passive, diode half-bridge rectifier (D_1, C_3, R_6, R_7) and a passive low-pass filter (R_8, C_4), both fed to the inputs of the comparator U_3 (TLV3701, Texas Instruments) which is the only active component in the channel detector.

The principle of operation of the whole feature extractor's signal chain is shown by TINA-TI SPICE simulation in Fig. 4. Test signal (V_{mic}, gray) consisted of mono-component (single frequency, 3.6 kHz) in-band chirp of the duration from 0.12 to 0.17 s, windowed by a Hamming window, of SNR of 6 dB.

The filter extracts in-band spectral content and amplifies it with gain of 2 (V_{filt}, green). The diode rectifier extracts its envelope (V_{rect}, blue), while a large time constant of the low pass filter provides the baseline signal V_{int}, making the detection invariant of filter's output signal level. Finally, the comparator's digital output V_{out} (purple) tracks the duration of the in-band signal chirp. The comparator's detection sensitivity can be adjusted by the ratio of the divider R_6, R_7.

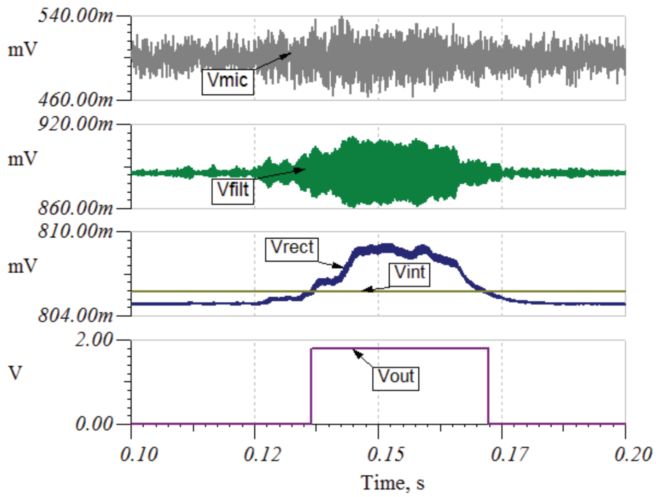

Fig. 4. SPICE simulation showing working principle of the implemented circuit. From top to bottom: microphone's input signal V_{mic} (gray), filter's output V_{filt} (green), rectified signal envelope V_{rect} (blue) and baseline V_{int}, channel detector's output V_{out} (purple).

V. Experimental Results

To evaluate the proposed architecture in terms of functionality and power consumption, we designed a modular prototype printed circuit board (PCB) implementing individual channels of the acoustic detector. It includes MEMS microphone (ICS-40310), the GIC filter using OPA2379 amplifiers, passive in-band signal duration detector with zero-bias HSMS282x Schottky diode, and TLV3701 comparator.

A power supply featuring linear voltage regulator (TPS79718) and reference (ISL21080) is included. Potentiometers connected in series with R_1, R_2, R_3 and R_5 enable programming of filter and tuning of detector respectively. A photograph of the prototype PCB is shown in Fig. 5. In the rest of the paper, we show experimental results obtained on the prototype board: (1) measurement of frequency characteristics of the filter, (2) verification of detector's response in multi-channel environment and (3) power consumption analysis.

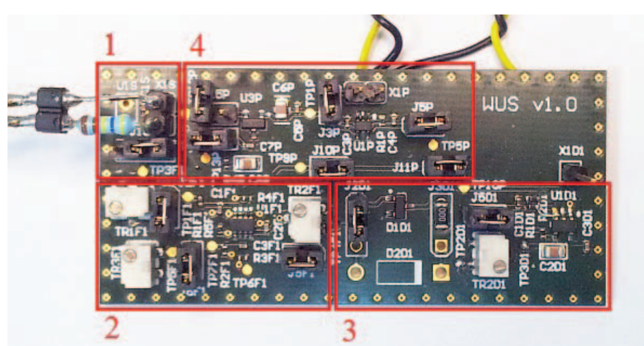

Fig. 5. Modular PCB prototype implementing single channel of the spectro-temporal feature extraction circuit. 1 – microphone or external signal input, 2 – tunable channel filter, 3 – channel detector, 4 – power supply.

A. Filter's frequency response measurement

Amplitude and phase characteristics of the filter were measured using signal analyzer HP35670A. The MEMS microphone input was simulated using a test signal of amplitude of 2 mV_{pp}, 500 mV DC offset and additional 4.5 kΩ external impedance. Characteristics were measured with filter's central frequency preset to {0.5, 1, 2, 3, 4} kHz, and bandwidth set to 200 Hz. Fig. 6 shows the obtained results (red line), comparing them to the SPICE simulation (blue). For equal preset bandwidth, measured results exhibit better selectivity (sharper roll-off). As seen, both the measured and the simulated amplitude characteristics exhibit a slight shift in the obtained central frequency, proportional to preset frequency (central frequency preset error) due to the limited gain-bandwidth product of the OPA2379 amplifier. The effect may be compensated by one-time calibration. Also, deviation from the expected central-frequency gain of 6dB can be observed, proportional with frequency.

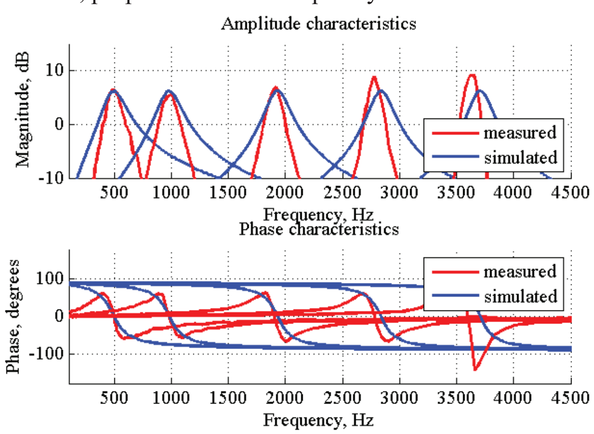

Fig. 6. Comparison of measured and simulated frequency response of the GIC filter. Up – amplitude, down – phase characteristics.

B. Detectors' multi-channel response verification

To verify the concept of multi-channel spectro-temporal wake-up audio pattern detection, a two-channel system was assembled from the prototype PCB modules described in Section III.C.

For test purposes, central frequencies of two channel filters were preset to 1.9 and 3.8 kHz respectively, and bandwidth of 200 Hz, to correspond to the example of used synthetic test signal, with spectro-temporal content shown in Fig. 7 (upper portion). NI-6211 DAQ and Matlab were used for test signals generation, DAC conversion, and capture of the detectors' responses.

As seen from lower part of Fig. 7, the signal's spectral content is correctly temporally localized by both channels. The detectors' selectivity is demonstrated by no spurious comparator activation throughout duration of the interference signal at 2.8 kHz (from 3rd to 4th second). The limit of the detectors' sensitivity, defined by occurrence of either false positively or false negatively classified temporal intervals, was with signals of amplitudes lower than 4 mV, buried in 2 mV_{pp} white noise (e.g. down to SNR = 6 dB).

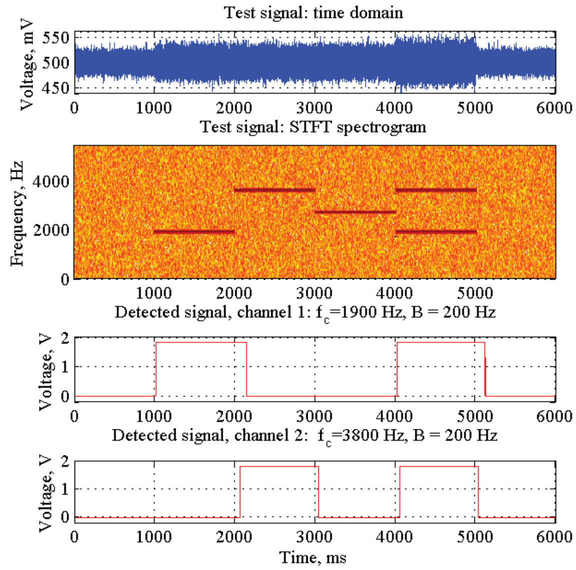

Fig. 7. Example of multichannel signal detection showing two in-band chirps of 1900 and 3800 Hz localized in time by two respectively tuned channel filters and detectors.

C. Power consumption analysis

Power consumption analysis is based upon measurements of average supply currents of the prototype PCB's subsystems: microphone, filter, detector, and power supply. Keithley 485 picoampermeter was used for measurement. Total supply average current of the PCB board measured at 2.5 V external DC supply voltage was 25.01 µA. Table 1 breaks power results down per individual subsystem, showing that total power required for the multi-channel system would be in order of (18.55 + N×11.41) µW, where N is number of channels. A rough estimate of the power consumption savings gained by the usage of our analog-domain feature extraction circuit, is shown by comparison with a digital signal processing chain implementing the analogous functionality. Chain consisted of the microphone ICS-40310 connected to a low-power amplifier (OPA433), and MCU MSP430FR5969 (operating at 3 V, 8 MHz) with integrated analog-to-digital converter (ADC). Signal was amplified, digitized at 8 Kb/s, and band-pass digital filtering on MCU. Under assumption of MCU operating in always-on mode, last row in Table 1 shows the significant increase in average total power consumption, in the range of 3 mW.

TABLE I. POWER CONSUMPTION PER SUBSYSTEMS

Subsystem	Current, µA	Voltage, V	Power, µW
Microphone	17.80	0.90	16.02
Filter	5.56	1.80	10.01
Detector	0.78	1.80	1.40
Power supply	1.01	2.50	2.53
Total power per channel			**29.96**
Total power, multi-channel (N = 2)			**41.37**
Total power, multi-channel (N = 3)			**52.78**
Digital implementation (MSP430, 3 V, 8 MHz)			**3240**

More detailed analysis of the filter's power consumption is provided in Fig. 8, showing its dependency on preset filtering parameters and properties of input signal. It can be seen that power lowers with a higher preset filtering frequency, and by lowering bandwidth, for 3 μW. As for the input signal's voltage, its AC magnitudes range is much lower (1-20 mV), in comparison to its DC offset of around 450-550 mV. Thus, DC offset dominates the power consumption, and may affect the filter's power by more than 1 μW.

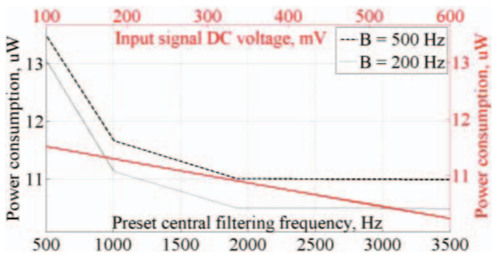

Fig. 8. Parametric plot of power consumption per channel with respect to preset filtering frequency and input signal DC magnitude.

VI. CONCLUSIONS

In context of energy-efficient acoustic monitoring, we proposed a low-power, always-on, audio pattern wake-up stage, keeping the main processing stage asleep in absence of potential events. Novel design of the spectro-temporal feature extraction part of the wake-up circuit was presented, consisting of multi-channel filter bank and channel detectors.

Filter characterization demonstrated the benefits of GIC filtering topology for low-power applications in the audio frequency range. The detector's response was experimentally tested in multi-channel environment, verifying its capability to detect audio signals as low as 4 mV in magnitude, at SNR of 6 dB. Power consumption, ranging from 10.5 to 13.5 μW per channel, was measured. Our results, obtained on prototype assembled with discrete off-the-shelf components, are comparable to state-of-the-art integrated circuit designs [14]-[16] (see Section II), and show potential for further reduction.

Our future efforts are aimed at bringing the design closer towards integrated audio wake-up circuit, by adding circuitry for dynamic tuning of channel filters' and detectors' parameters and power. Finally, the wake-up circuit shall be completed by the design of low-power classification circuit.

ACKNOWLEDGMENT

This work was supported by "Transient Computing Systems", SNF project (200021_157048), by SCOPES SNF project (IZ74Z0_160481), and by ETHZ Grant funding.

REFERENCES

[1] Puccinelli, D., & Haenggi, M. (2005). Wireless sensor networks: applications and challenges of ubiquitous sensing. Circuits and Systems Magazine, IEEE, 5(3), 19-31.

[2] Patel, S., Park, H., Bonato, P., Chan, L., & Rodgers, M. (2012). A review of wearable sensors and systems with application in rehabilitation. J Neuroeng Rehabil, 9(12), 1-17.

[3] Abbasi, A. Z., Islam, N., & Shaikh, Z. A. (2014). A review of wireless sensors and networks' applications in agriculture. Computer Standards & Interfaces,36(2), 263-270.

[4] Dargaville, T. R., Farrugia, B. L., Broadbent, J. A., Pace, S., Upton, Z., & Voelcker, N. H. (2013). Sensors and imaging for wound healing: a review.Biosensors and Bioelectronics, 41, 30-42.

[5] Dargie, W. (2012). Dynamic power management in wireless sensor networks: State-of-the-art. Sensors Journal, IEEE, 12(5), 1518-1528.

[6] Popovici, E., Magno, M., & Marinkovic, S. (2013, June). Power management techniques for wireless sensor networks: a review. In Advances in Sensors and Interfaces (IWASI), 2013 5th IEEE International Workshop on (pp. 194-198). IEEE. Chicago

[7] D. Oletic, B. Arsenal, V. Bilas, "Low-Power Wearable Respiratory Sound Sensing", MDPI Sensors, 2014, 14, pp. 6535-6566

[8] Lauwereins, S., Badami, K., Meert, W., & Verhelst, M. (2015). Optimal resource usage in ultra-low-power sensor interfaces through context-and resource-cost-aware machine learning. Neurocomputing.

[9] S. Lauwereins, K. Badami, W. Meert, M. Verhelst, Context- and cost-aware feature selection in ultra-low-power sensor interfaces, in: ESANN, 2014, pp. 93–98.

[10] Ramakrishnan, S.; Basu, A.; Leung Kin Chiu; Hasler, J.; Anderson, D.; Brink, S., "Speech Processing on a Reconfigurable Analog Platform," Very Large Scale Integration (VLSI) Systems, IEEE Transactions on , vol.22, no.2, pp.430,433, Feb. 2014

[11] M. Malinowski, M. Moskwa, M. Feldmeier, M. Laibowitz, and J. A. Paradiso, "Cargonet: a low-cost micropower sensor node exploiting quasipassive wakeup for adaptive asychronous monitoring of exceptionalevents," in Proceedings of the 5th international conference on Embedded networked sensor systems. ACM, 2007, pp. 145–159.

[12] S. Jevtic, M. Kotowsky, R. P. Dick, P. A. Dinda, and C. Dowding, "Lucid dreaming: reliable analog event detection for energy-constrained applications," in Proceedings of the 6th international conference on Information processing in sensor networks. ACM, 2007, pp. 350–359.

[13] T. Delbruck, T. Koch, R. Berner, and H. Hermansky, "Fully integrated 500uw speech detection wake-up circuit," in Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on. IEEE, 2010, pp. 2015–2018.

[14] A. G. Katsiamis, E. Drakakis, and R. F. Lyon, "A biomimetic, 4.5 uW, 120+ db, log-domain cochlea channel with AGC," Solid-State Circuits, IEEE Journal of, vol. 44, no. 3, pp. 1006–1022, 2009.

[15] Rumberg, B.; Graham, D.W., "A Low-Power and High-Precision Programmable Analog Filter Bank," Circuits and Systems II: Express Briefs, IEEE Transactions on , vol.59, no.4, pp.234,238, April 2012

[16] Rumberg, Brandon, and David W. Graham. "A low-power magnitude detector for analysis of transient-rich signals." Solid-State Circuits, IEEE Journal of 47.3 (2012): 676-685.

[17] R. Sarpeshkar, C. Salthouse, J.-J. Sit, M. W. Baker, S. M. Zhak, T. K.-T. Lu, L. Turicchia, and S. Balster, "An ultra-low-power programmable analog bionic ear processor," Biomedical Engineering, IEEE Transactions on, vol. 52, no. 4, pp. 711–727, 2005.

[18] Yamasaki, Toshihiko, and Tadashi Shibata. "Analog soft-pattern-matching classifier using floating-gate MOS technology." Neural Networks, IEEE Transactions on 14.5 (2003): 1257-1265

[19] Lazzaro, John, John Wawrzynek, and Richard P. Lippmann. "A micropower analog circuit implementation of hidden Markov model state decoding." Solid-State Circuits, IEEE Journal of 32.8 (1997): 1200-1209.

[20] W.-K. Chen, The circuits and filters handbook. CRC Press, 2002.

[21] OPA2379 1.8V, 2.9μA, 90kHz, Rail-to-Rail I/O Operational Amplifiers, Datasheet SBOS347D, Texas Instruments, 2008

[22] ICS-40310 Ultra-low Current, Low-Noise Microphone with Analog Output Datasheet, Invensense, 2014

[23] M. Magno and L. Benini, "An ultra low-power high sensitivity wakeup radio receiver with addressing capability," in Wireless and Mobile Computing, Networking and Communications (WiMob), 2014 IEEE 10th International Conference on. IEEE, 2014, pp. 92–99.