A LOW-POWER INTEGRATED CIRCUIT FOR ANALOG SPIKE DETECTION AND SORTING IN NEURAL PROSTHESIS SYSTEMS

A. Bonfanti[†§], T. Borghi[§], R. Gusmeroli[§], G. Zambra[§], A. S. Spinelli[§], A. Oliynyk[*]
L. Fadiga[†*] and G. Baranauskas [†]
[†]*Department of Robotics, Brain and Cognitive Sciences, The Italian Institute of Technology - Genova, Italy*
[§]*Dipartimento di Elettronica e Informazione, Politecnico di Milano - IU.NET, Milano, Italy*
[*]*Dipartimento di Scienze Biomediche e Terapie Avanzate, Sezione di Fisiologia umana, Universitá di Ferrara - Ferrara, Italy*

Keywords: Prosthetic device, Action potential (AP), Multichannel recording system, Spike detection, Spike sorting.

Abstract: Since the proof that prosthetic devices directly controlled by neurons are viable, there is a huge increase in the interest in integrated multichannel recording systems that register neural signals with implanted chronic electrodes. One of the bottlenecks in such compact systems is the limited rate of data transmission in the wireless link, requiring some sort of data compression/reduction. To solve such a problem, we propose an analog low power integrated system for action potential (AP) detection and sorting. In this system, AP detection is performed by a double threshold method that reduces the probability of false detections while AP sorting is based on the measurement of peak and trough amplitudes and spike width. The circuit has been implemented in $0.35-\mu$m CMOS technology with power consumption of 70 μW per channel including the pre-amplifier. The system was tested with real traces. Compared to standard AP sorting techniques, the proposed simple AP sorter was able to correctly assign to single units over 90% of detected APs. Thus, our system preserves most information encoded by APs and we estimate that for a typical trace the required bandwidth per channel will be less than 4 kbps or 400 kbps for 100 channel.

1 INTRODUCTION

In recent years, in a number of laboratories it has been shown that signals obtained with multichannel extracellular recordings are sufficient to control the movements of a prosthetic device in real time (Hochberg et al., 2006). However, we are still far away from any device to be clinically tested because of enormous technological and scientific challenges (Stieglitz, 2007). According to current thinking, a future neural prosthetic device should contain a set of neural amplifiers connected to a signal-processing unit and a wireless link to transmit the amplified signals down to an actuator (*e.g.*, a robotic arm, a remote controller, a mouse, and so on). Unfortunately, the low power state-of-the-art wireless systems cannot transmit data at the high rate required for this kind of application where power budget is limited to 800 μW/mm^2 (Harrison and Charles, 2003). For example, 100 channels sampled at 30 kHz with 10 bits resolution would require to transmit 30 Mbps while low power wireless systems can handle only $<$ 3 Mbps (Olsson et al., 2005). Thus, to enable data transmissions from multichannel neural amplifiers, it is necessary to compress or reduce raw neural data. Since in a typical neural trace only action potentials (APs) contain any useful information, a very efficient way of neural signal reduction is to transmit a single bit for each AP. Researchers in (Harrison et al., 2007) used simple threshold detectors to identify APs in the amplifier signal and then transmitted zeros when no AP were detected and ones for each threshold crossing event. However, each recording site usually collects the activity from several nearby neurons and the aforementioned method does not provide information on which neuron fired at a given time. This may be a serious drawback in motor prostheses applications where it is important to isolate the activity of each neuron to better predict the intended movement (Hochberg et al., 2006). This paper presents a low-power circuit that permits neuron discrimination while drastically reducing the required data bandwidth per channel. This circuit extracts three parameters for a single spike, namely the

Bonfanti A., Borghi T., Gusmeroli R., Zambra G., S. Spinelli A., Oliynyk A., Fadiga L. and Baranauskas G. (2009).
A LOW-POWER INTEGRATED CIRCUIT FOR ANALOG SPIKE DETECTION AND SORTING IN NEURAL PROSTHESIS SYSTEMS.
In *Proceedings of the International Conference on Biomedical Electronics and Devices*, pages 67-74
DOI: 10.5220/0001539200670074

peak and trough amplitudes and the time between the peak and the trough. It has been shown that the use of these AP features is sufficient to adequately sorts APs and is actually superior to spike spike separation using more parameters (Vibert and Costa, 1979). According to (Olsson et al., 2005), for sampling rate of 20 kHz at 5 bit resolution, a similar AP feature extraction method yields data compression close to 90%. Although a similar in function analog device has been described (Horiuchi et al., 2004) a digital extraction of these three parameters has been presented also (Olsson et al., 2005), in this work we introduce significant innovations: a) the implementation of a new spike-detection algorithm that decreases several-fold the rate of false detections; b) an improved peak and trough detectors with very fast recovery time permitting the detection of closely spaced APs, and c) the use of an additional feature, namely the time difference between peak and trough occurrence. Moreover, the amplification of neural signals is performed with an amplification topology that provides the best known trade-off between noise and power consumption. To confirm the efficiency of our device, we tested it employing artificial and real signals and we conclude that, in the majority of cases, with few independent units present in a trace, the extracted three parameters can be sufficient to achieve adequate single unit separation.

2 SYSTEM ARCHITECTURE

The system architecture is shown in Fig. 1 and comprises a low-noise amplifier (LNA), a spike detector and a spike sorter. The spike detector splits the amplified signal in two paths and exploits a double-threshold strategy, which will be discussed in detail in Section 2.2. The signal coming from the LNA is also processed by the spike sorter, which extracts from each detected spike three features, namely the peak and trough amplitudes and the time interval between peak and trough. The detailed circuit and its performance descriptions follows below.

2.1 Low-power and Low-noise Neural Amplifier

Pre-amplification is obtained through a double-stage active filter, shown in Fig. 2. The first stage is an ac-coupled inverting high-pass filter, employing two MOS-bipolar pseudoresistors as feedback elements (Harrison and Charles, 2003). This enables the synthesis of high-value resistance without large area consumption. Mid-band gain is set to $G_1 = -C_1/C_2 = -66$, with $C_1 = 10$ pF and $C_2 = 150$ fF. The higher cut-off frequency is determined by the gain-bandwidth product of the first operational amplifier ($GBWP_1$) and is set to about $GBWP_1/G_1 = 20$ kHz, while the high-pass pole frequency was designed to be about 10 Hz as determined by C2 and MOS pseudoresistors . The second stage was designed to reject the first stage offset while providing further signal amplification and filtering. This stage employs the same pseudoresistor elements in the feedback path, with a gain $G_2 = -50$ (achieved with $C_3 = 7.5$ pF and $C_4 = 150$ fF). The operational amplifier in the first stage was designed to have an input-referred noise in the pre-amplifier band lower than 5 μV_{rms}. Furthermore, since this circuit is thought to be used in an implantable system with a large number of electrodes, the current consumption was minimized by designing a telescopic operational amplifier (see Fig. 3); in fact, this well-known configuration gives excellent noise performance thanks to the limited number of transistors and can achieve the best trade-off between power consumption and noise. The input PMOS and current mirror NMOS transistors are biased in weak-inversion and strong inversion region (Harrison and Charles, 2003), respectively, since $(W/L)_p \gg (W/L)_n$ (see Table 1, where IC is the inversion coefficient which is $<< 1$ in the

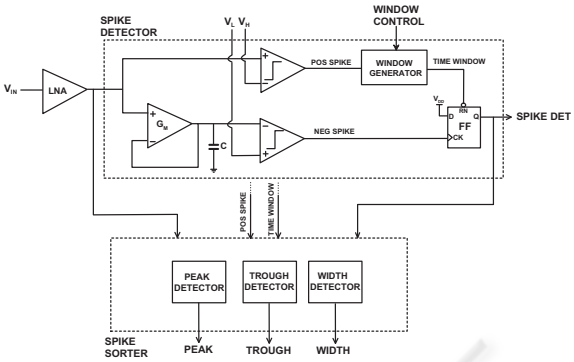

Figure 1: Block scheme of the detection and sorting system.

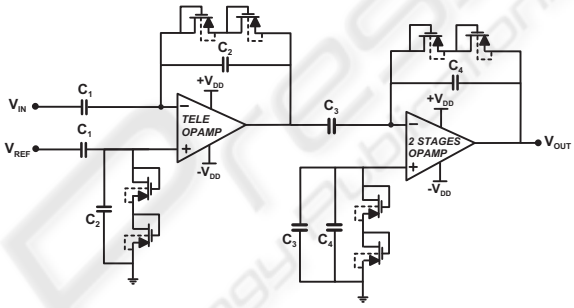

Figure 2: Overall schematic of the pre-amplifier (LNA).

weak-inversion region). This implies that M_P transistors have a transconductance much higher than M_N ones and hence the equivalent input-referred noise is mainly determined by the the input transistors of the differential pair. Considering both the flicker noise and the thermal noise, the input referred power noise density results:

$$\overline{E_{n_{eq}}^2} = \frac{8KT\gamma}{g_{mp}} + \frac{2K_p^{(1/f)}}{C'_{ox}W_pL_p}\frac{1}{f}, \quad (1)$$

with γ equal to $1/(2\kappa)$ in weak-inversion region ($\kappa \approx 0.7$). In this working region the transconductance may be approximated by:

$$g_m \cong \frac{\kappa I_D}{U_T}, \quad (2)$$

and so the input-referred thermal noise spectral density results:

$$\overline{E_{n_{th}}^2} = \frac{8kTU_T}{\kappa^2 I_{bias}}, \quad (3)$$

where I_{bias} is the total current consumption of the telescopic cascode amplifier. Setting an upper limit for the thermal noise contribution in the 20-kHz amplifier band of 3 μV_{rms}, a bias current of 4 μA has to be chosen. The flicker noise contribution was carefully minimized choosing a large area for the input PMOS transistors, $W_p \times L_p = 400 \times 1\ \mu m^2$, that implies a noise corner frequency lower than 100 Hz assuming a flicker noise coefficient $K_p^{(1/f)} \approx 2 \times 10^{-26}$ V^{2}F. A figure of merit proposed in (Harrison and Charles, 2003) and widely adopted to compare different design is the Noise Efficiency Factor (Harrison and Charles, 2003), defined as:

$$NEF = V_{in,rms}\sqrt{\frac{2I_{tot}}{\pi U_T\, 4kT\, BW}}, \quad (4)$$

where $V_{in,rms}$ is the input-referred rms noise, I_{tot} is the total supply current and BW is the bandwidth of the amplifier. For the aforementioned sizing strategy, one can find from Eqs. (4) and (3) that the minimum NEF value is $\sqrt{2}/\kappa = 2.02$, which is the minimum value achieved with a conventional differential pair as input stage (Wattanapanitch et al., 2007).

Table 1: Transistor Operating Points.

	(W/L) μm	V_{ov}[mV]	IC	g_m[μA/V]
M_P	400/1	103	0.17	70
M_N	50/1	535	61	8
M_{CAS}	5/40	17	0.56	50

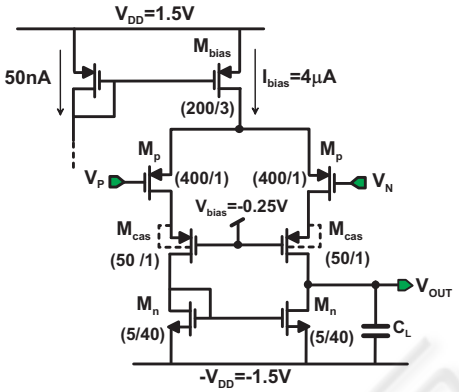

Figure 3: Schematic of the telescopic operationaL amplifier used in the pre-amplifier first stage.

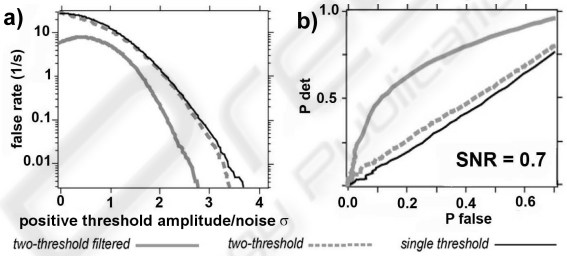

Figure 4: Performance of different spike detection algorithms. a) False detection rate as a function of the threshold amplitude normalized by σ; b) Comparison of receiver operating characteristics (ROC) in a low-SNR condition. P_{det} and P_{false} represent the probabilities to correctly and falsely detect a spike.

2.2 Spike-detection Algorithm and Circuit Implementation

The amplified signal is then sent to a spike-detection circuit. The detection algorithm is based on simultaneous detection of the peak and trough of a spike using two different thresholds, with a low-pass filter placed on the trough detection channel. A spike is detected only if a trough is found within a time interval after the peak. For simplicity, we will call the proposed algorithm that includes low-pass filtering for trough detection the *two-threshold with filtering* method, while the same algorithm without filtering will be simply called the *two-threshold* method. MATLAB® simulations were first run to assess the performance of the algorithm. For tests, the real-spike waveforms obtained from *in-vivo* and *in-vitro* recordings were used to build the simulated traces by superimposing a series of non-overlapping APs evenly spaced at 5.5 ms onto a realistic noise (taken from a recording without any spikes). The overall performance of the algorithm was evaluated both by determining for different threshold values the probability

of false detections alone and by the receiver operating characteristic (ROC). Fig. 4a shows that, for low thresholds, the two-threshold with filtering method can reduce two- to threefold the rate of false detections compared to a standard single-threshold method, while the improvement is less than 20% if no filter is used. Such a difference is explained by the fact that the trough is usually slower and smaller than the peak, and, without a proper filtering, the trough detector is much more likely to respond to the background noise. Meanwhile, ROC analysis Fig. 4b shows that the two-threshold with filtering method performs much better than the other two methods, especially for low signal-to-noise ratio (SNR). For instance, for a SNR (defined as the ratio between the peak-to-trough amplitude and six times the standard deviation (σ) of the noise) of 0.7, this method still detects APs reasonably well while the other two methods perform no better than detection by chance. Based on these results, we designed an analog circuit to implement the two-threshold with filtering algorithm. As shown in Fig. 1, the conditioned signal is split in two paths. The first signal path has a comparator with a selectable positive threshold V_H. Whenever the threshold is crossed, the comparator triggers a time-window generator, depicted in Fig. 5. The output of the set-reset flip-flop is normally in a low logic state $(Q = 0)$, keeping the capacitor charged. When a positive peak of AP is detected, the NMOS swiftly discharges the capacitor and the resulting drop of the transistor drain voltage turns off the same NMOS, thus enabling the capacitor re-charging. When the recharging capacitor voltage crosses the trigger threshold, the cycle is repeated. This block acts then as a clock generator, feeding the counter that generates the output, namely the time window. When the counter completes its cycle, it resets the flip-flop, bringing everything to the initial state. The current I_{CHARGE} can be externally tuned in order to produce time windows of different lengths. For example, by setting $I_{CHARGE} = 20$ nA we obtain a time window of 1 ms with a 10 pF capacitor. The second signal path has a low-pass filter (implemented with a $G_M - C$ cell with tunable bandwidth) followed by a second comparator with negative threshold V_L. Finally, outputs of the two channels feed a D-type flip-flop acting as an AND port: its output is true only when a negative crossing happens during the assertion of the time window. The following parameters of the circuit were used for the tests here: 5 kHz for the low-pass filter, 0.6 for the ratio of the two threshold amplitudes and about 2 ms for the time window duration.

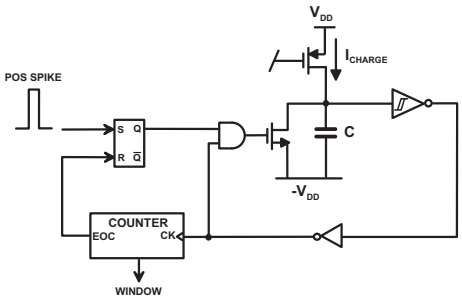

Figure 5: Simplified schematic of the time-window generator.

2.3 Spike-sorter Circuit

The amplified signal is also processed by a spike-sorter block (see Fig. 1) that measures the peak and the trough amplitudes and the width of the detected spike. The peak detector circuit is shown in Fig 6(a). If the *ENABLE* signal is low, the circuit acts as a traditional voltage follower, while when *ENABLE* is high the source follower M_6 is biased with the small leakage current of transistor M_7 and it is able to follow only the rising edge of the input signal. In the implemented circuit, the *ENABLE* signal comes from the time-window generator output: when the amplified signal crosses the positive threshold, the time-window signal is set high, forcing the detector to track the peak. This control signal allows detecting peaks with very different amplitudes and very close to each other, that was not possible with the implementation in (Horiuchi et al., 2004) because of the slow discharge after a peak detection. Two aspects were taken into consideration in the design of the peak detector. First, the $M_1 - M_4$ amplifier gain was maximized (to about 50 dB) in order to reduce the detector offset (which could cause a systematic error in the peak amplitude evaluation). Second, the gate-source capacitance of the follower transistor M_6 was minimized in order to keep low the voltage drop caused by the input voltage that drops below the peak voltage.

The trough detector is the PMOS equivalent of the Fig 6(a) circuit. In this case, the enable signal for the trough detector is the spike-detector output itself.

The width detector circuit, shown in Fig. 6(b), is a time-to-amplitude converter, based on the charging of a capacitor with a constant current, with two control signals, *ENABLE* and $\overline{RESET}$. The latter is the time-window signal (*TIME WINDOW*), while *ENABLE* is the logic exclusive OR between *TIME WINDOW* and *SPIKE DET*. The crossing of the positive threshold activates the current generator, M_5, and switches off M_6. The capacitor $C = 10$ pF is charged by a 30 nA constant current until the crossing of the negative

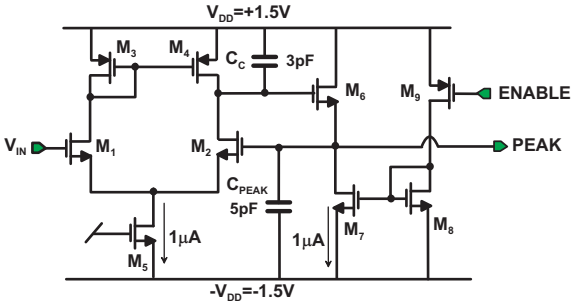

(a) Peak detector

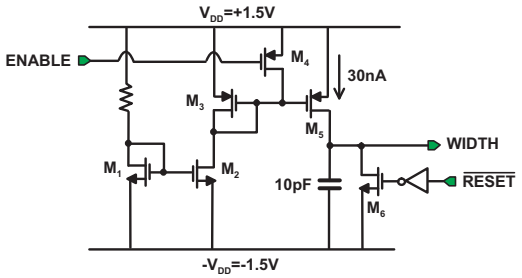

(b) Width detector

Figure 6: Simplified schematic of the peak detector (a) and width detector (b).

threshold. Then, the capacitor voltage is kept constant to a value proportional to the time between the positive and the negative threshold crossing, until the time window signal goes down discharging the capacitor. The gain of the width detector is about 3 V/ms. Note that all the three sorter output signals (peak, trough and width signals in Fig. 1) are meaningful only when the spike detector output is high. If a spike is detected, both the falling edge of the *TIME WINDOW* signal and of the *SPIKE DET* signal reset the three spike sorter outputs meaning that the *SPIKE DET* signal may be used as a trigger for read-out of the analog sorter outputs in order to convert these signals into a digital format.

3 CIRCUIT TEST

A test structure with a pre-amplifier, a spike detection circuit and an analog sorter was implemented in 0.35-μm AMS CMOS process occupying a total area of 0.24 mm^2 (see Fig. 7). The LNA was characterized using an HP35665A Dynamic Signal Analyzer. The measured transfer function is presented in Fig. 8, showing a mid-band gain of 71.5 dB and a 29 Hz$-$22 kHz pass-band. The same figure shows the transfer function of the LNA followed by the $G_M - C$ filter present in the trough path of the spike detector. The upper cut-off frequency is set to about 5 kHz by tuning the bias current of the $G_M - C$ low-pass cell. The measured input-referred noise power spectrum is shown in Fig. 9. It was obtained by dividing the output noise power spectral density by the square of the overall transfer function. The input-referred thermal noise is about 25 nV/$\sqrt{Hz}$, as expected for the set bias current, while the noise corner frequency is about 800 Hz, higher than the estimated value. This result is likely to be due the incorrect model of the flicker noise for a transistor working in the weak-inversion/subthreshold region, that is the case for the input PMOS transistors. Integrating the input-referred noise spectrum over the pre-amplifier band, we get a 4 μV_{rms} input noise, while the measured current consumption is 4.7 μA for the first stage LNA and 0.5 μA for the operational amplifier in the second stage. The Noise Equivalent Factor (NEF) of the LNA is 2.45 which is the best result reported to date (Fig. 10). However, this NEF is higher than the theoretical limit obtained in the previous section, due to the presence of the unaccounted flicker noise and of a second amplification/filtering stage. The spike detection block was tested with simulated APs signals immersed in a realistic background noise fed into the integrated circuit using a TTI TGA12104 arbitrary waveform generator. Fig. 11 shows the amplified artificial neural activity, the control signals that trigger the spike detector (output of positive- and negative-threshold comparators) and the spike detector output itself. Note the false-positive threshold crossing shown in figure: a simple threshold algorithm would interpret it as a true spike firing event, while our algorithm rejects it correctly, thus increasing the detection reliability. The current consumption of the spike detector is largely determined by the comparators (each one draws 2 μA), and by the time window generator. The latter dissipates only when the positive threshold is crossed and in the case of 100 Hz firing rate, the average current consumption is 10 μA. The $G_M - C$ filter has a negligible current consumption since its bias current is about 0.1 μA. In addition, to verify the correct measurement of the three features (peak, trough and width amplitudes) for detected spikes, the spike sorter circuit was tested with simulated waveforms. The amplified trace was recorded as well as peak, trough and width outputs of the spike sorter. A snapshot of the measured signals is shown in Fig. 12. The three features measured by the spike sorter were compared with the same features extracted from the acquired LNA output using MATLAB®. For amplified spikes of 400 mV peak amplitude the error was never larger than 30 mV. This value corresponds to an error of 8 μV for a 106 μ V peak amplitude input

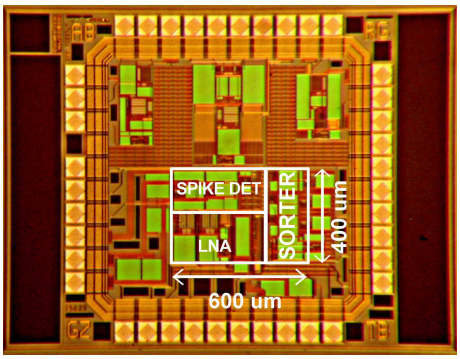

Figure 7: Die photo of the proposed circuit.

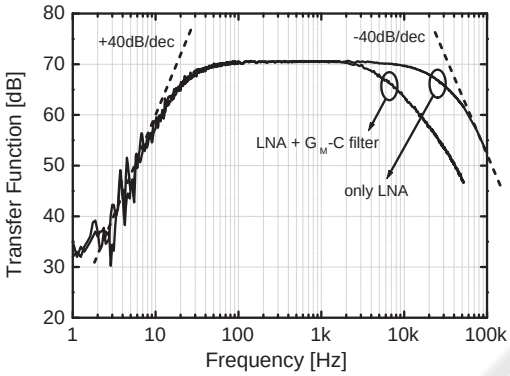

Figure 8: Measured transfer functions for the LNA amplifier (22-kHz bandwidth) and for the LNA + $G_M - C$ filter (5-kHz bandwidth).

spike. In the same way we obtained a maximum error of 50 μs for a spike width of 500 μs. The power consumption of the spike sorter block was kept very low: the peak and the trough detector draw 2 μA each while the width detector has a negligible consumption. The power dissipation of the whole system (output buffers excluded) is about 70 μW fora 3-V power supply. The main characteristics of the whole circuit are summarized in Table 2.

4 TEST WITH SIGNALS

To test the efficiency of the spike sorter, we used a variety of traces obtained during *in-vivo* recordings in monkeys employing tungsten electrodes (impedance in the range of $0.5 - 2$ MΩ). The signals were downloaded to the waveform generator and then fed into the IC through an attenuation filter to simulate the real amplitude of the neural signal and the input resistance of typical electrodes. For AP detection and sorting, the peak threshold was set to about 4 times the standard deviation of the background noise while

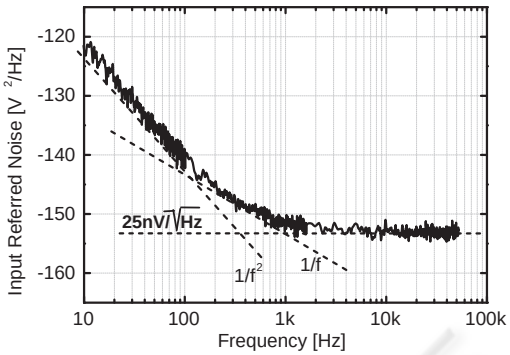

Figure 9: Measured input-referred noise spectrum of the LNA.

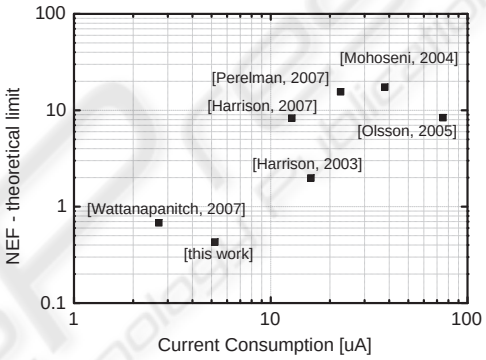

Figure 10: Difference between measured *NEF* and minimum theoretical limit (2.02) of published neural differential amplifiers as a function of the supply current.

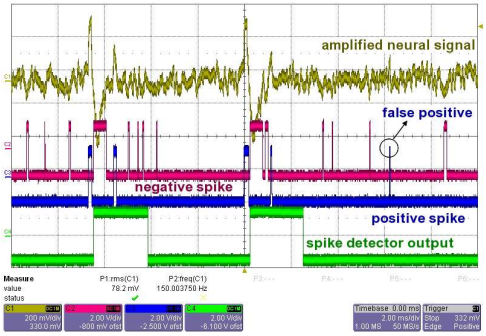

Figure 11: Snapshot of measured signals. From top to bottom: LNA output signal, outputs of the comparator with negative and positive threshold and spike detector output.

the trough detector threshold was 0.6 times the peak threshold. All these tests yielded similar results, so that here we only describe a typical case obtained with an *in vivo* recording trace (see Fig. 13). We compared the chip sorted waveforms with the results obtained employing the Principal Components (PCs) analysis of AP waveforms and identifying clusters of events in the space of the three largest PCs (Lewicki, 1998).

Table 2: System electrical characteristics.

Technology	CMOS 0.35−μm AMS
Amplifier + spike-detector + spike sorting area	0.24 mm^2
Power supply	3.0 V
Pre-amplifier midband gain	71.5 dB
Pre-amplifier bandwidth	29 Hz−22 kHz
Pre-amplifier input noise (29 Hz−22 kHz)	4 μV_{rms}
Overall pre-amplifier NEF	2.45
Current consumption (output buffers excluded)	
Pre-amplifier (1^{st} stage)	4.7 μA
Pre-amplifier (2^{nd} stage)	0.5 μA
Spike detector (for a 100 Hz firing rate)	14.1 μA
Spike sorter	4.1 μA
Total	23.4 μA

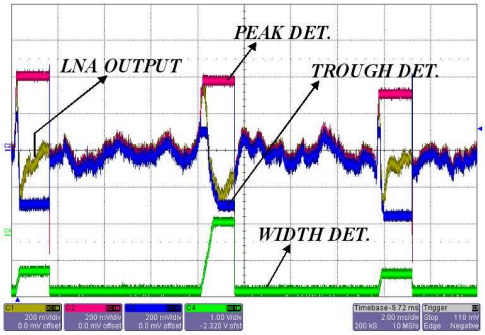

Figure 12: Snapshot of the measured signals of the analog sorter: amplified neural signal (gold line), peak detector output (red line), trough detector output (blue line) and width detector output (green line).

In this case, the AP waveforms were obtained by detecting the events of threshold crossing and collecting about 0.3 ms of trace before and 1.3 ms after the event. Since we used real data, there is no assurance that AP shapes are sorted correctly. Therefore, to asses the quality of unit separation, we used AP occurrence auto- and cross- correlograms (Harris et al., 2000). The method is based on the fact that each neuron has a certain refractory period of $2-3$ ms after firing an AP during which no spike can be generated. Thus, if all AP shapes are derived from a single neuron, there should be no APs occurring at intervals less than the refractory period of the neuron. This absence of APs at brief intervals will correspond to no events around the midpoint (0 ms interval) in the auto-correlogram (Fig. 13). In the case of a pure noise signal, the probability to detect events around the midpoint (i.e.: events separated by small time intervals) would be equal to the probability of detecting events separated by large time intervals. Thus, by comparing the frequency of events at very short and very long intervals it is possible to estimate the quality of the achieved unit separation. In the example shown in Fig. 13, both methods were able to detect two units and the remaining events turned out to be noise. While in the PCA analysis both units have clean refractory periods with no events for intervals of < 3 ms, the chip separated units have a small background noise (Fig. 13). However, compared to the large interspike intervals, the frequency of these short interspike interval events was very low (< 10 fold) suggesting that less than 10% of identified units were misclassified. Similarly, the number of events detected by our spike sorter was $> 90\%$ of the ones detected by PCA analysis confirming that our simple sorter is able to achieve 90% efficiency of the PCA analysis method. Similar results were obtained with the simulated test signals downloaded from the public database (`http://www.vis.caltech.edu/~rodri/Wave_clus/Waveclus_home.htm`, data not shown) thus confirming the conclusions reached with the real signals.

5 CONCLUSIONS

The described analog integrated circuit for on-line AP waveform separation was designed with intention to be a part of a compact integrated multichannel system where low power consumption, small size and data compression/reduction are key factors to be considered during the systems design. The bandwidth required to transmit the three AP features depends on the neurons firing rate. For a typical electrode signal containing $2-3$ neurons firing rate at < 50 Hz, and with 10 bits of resolution, the estimated bandwidth to transmit one channel data processed by our device or less than $4-5$ kbps. Thus, information contained in 100 channels can be wireless transmitted with a device capable of 500 kbps data rates that is well within the range of current state of the art systems. The three features selected to perform the discrimination of different spikes were proved to work with real signals. Our implementation overcomes some of the problems

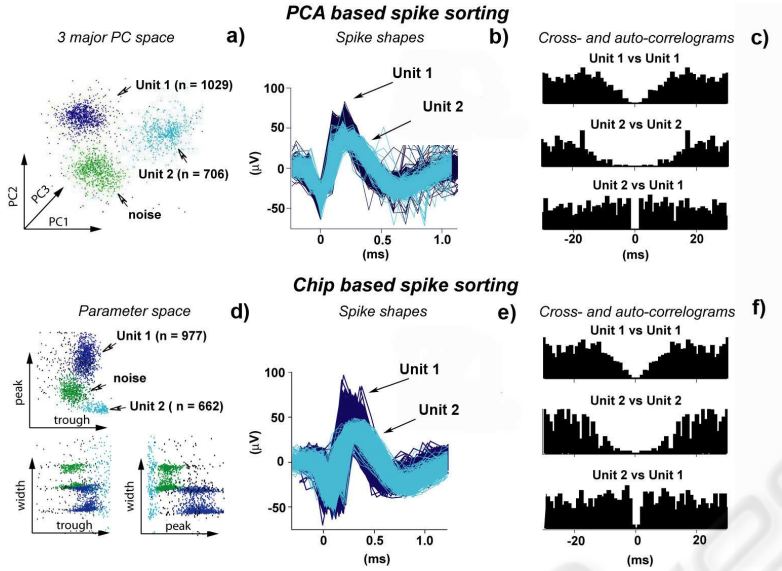

Figure 13: Comparison of PCA- and chip-based spike sorting. a) Unit 1 and 2 and noise plotted in the PCA space. b)Waveforms of single units. c) Cross- and auto-correlograms of waveforms selected with PC analysis. d) Single waveforms plotted in the chip-extracted three dimension space of features. Isolated clusters are clearly visible. e) Single units' waveforms assigned by employing the chip-based sorting. f) Cross- and auto-correlograms of the single units identified with the chip-based sorting method.

of similar systems, such as a slow recovery of peak detectors after sensing an AP (Horiuchi et al., 2004), while remaining an extremely low-power circuit with overall power density of ($< 300\ \mu W/mm^2$) that makes this circuit suitable for multichannel implantable systems.

ACKNOWLEDGMENTS

The authors would like to thank M. L. Grossi for her help in the circuit design. This work was partially supported by IIT (Italian Institute of Technology) and by EC funds to LF (RobotCub IP).

REFERENCES

Harris, K., Henze, D., Csicsvari, J., Hirase, H., and Buzsaki, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *Journal of Neurophysiology*, 84:401–414.

Harrison, R. and Charles, C. (2003). A low-power low-noise cmos amplifier for neural recording applications. *IEEE J. Solid-State Circuit*, (6):958–965.

Harrison, R. R., Watkins, P. T., Kier, R., Lovejoy, R., Black, D. J., Greger, B., and Solzbacher, F. (2007). A low-power integrated circuit for a wireless 100-electrode neural recording system. *IEEE J. Solid-State Circuit*, (1):123–133.

Hochberg, L., Serruya, M., Friehs, G., Mukand, J., Saleh, M., Caplan, A., Branner, A., Chen, D., Penn, R., and Donoghue, J. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, (442):164–171.

Horiuchi, T., Swindell, T., Sander, D., and Abshier, P. (2004). A low-power CMOS neural amplifier with amplitude measurements for spike sorting. In *Proceedings of ISCAS*, pages 2932–2395.

Lewicki, M. (1998). A review of methods for spike sorting: the detection and classification of neural action potentials. *Network: Computation in Neural Systems*, 9(4):53–78.

Olsson, R. and Wise, K. (2005). A three-dimensional neural recording microsystem with implantable data compression circuitry. *IEEE J. Solid-State Circuit*, (40):2796–2804.

Olsson, R. H., Buhl, D. L., Sirota, A. M., Buzsaki, G., and Wise, K. D. (2005). Band-tunable and multiplexed integrated circuits for simultaneous recording and stimulation with microelectrode arrays. *IEEE Trans. Biomed. Eng.*, 52:1303–1311.

Stieglitz, T. (2007). Neural prostheses in clinical practice: biomedical microsystems in neurological rehabilitation. *Acta Neurochir Suppl*, 97:411–418.

Vibert, J. and Costa, J. (1979). Spike separation in multiunit records: a multivariate analysis of spike descriptive parameters. *Electroenceophal. Clin. Neurophys.*, (47):172–182.

Wattanapanitch, W., Fee, M., and Sarpeshkar, R. (2007). An energy-efficient micropower neural recording amplifier. *IEEE Trans. Biomed. Eng.*, 1(2):136–147.