An implementation of a wavelet-based seizure detection filter suitable for realtime closed-loop epileptic seizure suppression

M.N. van Dongen, A. Karapatis, L. Kros, O.H.J. Eelkman Rooda, R.M. Seepers, C. Strydis, C.I. De Zeeuw, F.E. Hoebeek and W.A. Serdijn

Presented at the 10th IEEE Biomedical Circuits and Systems Conference, BioCAS 2014. DOI:10.1109/BioCAS.2014.6981773

ABSTRACT

This paper presents the design and implementation of a real-time epilepsy detection filter that is suitable for closed-loop seizure suppression. The design aims to minimize the detection delay, while a reasonable average detection rate is obtained. The filter is based on a complex Morlet wavelet and uses an adaptive thresholding strategy for the seizure discrimination. This relatively simple configuration allows the algorithm to run on a cheap and readily available microprocessor prototyping platform. The performance of the filter is verified using both in vivo real-time measurements as well as simulations over a pre-recorded EEG dataset (29.75 hours with 1914 seizures). An average detection delay of 492 ms is achieved with a sensitivity of 96.03% and a specificity of 93.60%.

4.1. INTRODUCTION

Several studies have shown the feasibility of a closed-loop system for epilepsy suppression by using electrical or optogenetic stimulation of various parts of the central nervous system[1–3]. These systems apply stimulation when ictal activity is detected in electroencephalogram (EEG) or electrocorticography (ECoG) recordings. Such an epilepsy detection system should detect a seizure as fast as possible in order to minimize the ictal time for the subject. Wavelets have been previously proposed as a tool to detect epileptic seizures, since they make it possible to look at the frequency and time properties of the signal simultaneously[4]. This paper presents the design of a wavelet based detection filter that aims to minimize the detection delay. The wavelet is implemented by a Finite Impulse Response (FIR) filter of which the impulse response corresponds to a truncated wavelet at a single scale. The detection is made with an adaptive thresholding mechanism. Due to the low complexity, it is possible to implement it using cheap and readily available off the shelve components, allowing for rapid prototyping, maximum flexibility and easy reproducibility. A realization using the credit-card sized Beaglebone microprocessor platform is discussed and the performance is analyzed and compared with other detection methods. Furthermore, *in vivo* measurements confirm the applicability of our system in a closed-loop setup.

4.2. BACKGROUND

The ECoG recordings used for this study are obtained from awake *Cacna1a*P601L (*tottering; tg*) mutant mice, that exhibit absence seizures. For the ECoG recordings Teflon coated silver ball tip electrodes (diameter 0.2mm, Advent Research Materials, Eynsham, Oxford, UK) are mounted bilaterally above the primary motor cortices and primary sensory cortices. The recorded signals from the head-fixed animals are amplified using a CyberAmp amplifier (Molecular Devices LLC., Axon Instruments, Sunnyvale, CA, USA). An epileptic absence seizure is characterized by massive synchronous discharging of thalamo-cortical neurons at a certain frequency 5Hz < f_e <8Hz. In the ECoG this translates into cortical generalized spike-and-wave discharges (GSWDs) that repeat themselves at this frequency. An epileptic seizure is defined as a continuous bursting of cortical GSWDs in the EEG for at least 1 s. Ictal activity should therefore be detected within a few hundred ms such that it can be suppressed before it develops into a seizure. A dataset consisting of 29.75 hours (split in traces of 15 minutes) of prerecorded ECoG data from 24 subjects was available to analyze the performance of the filter. Based on the above definition of an epileptic seizure, the data was annotated to determine time stamps at which seizures (N = 1914) start and end. These time stamps were obtained by an offline

algorithm that uses peak detection to find the GSWDs and labels an interval as a seizure when the GSWDs occur with a repetition rate of at least 5Hz during at least 1 s. The outcome of the algorithm was validated visually by a neuroscientist. Most studies on closed-loop seizure suppression give little details about the performance of the detection method used. In[3] simple thresholding is used, which is fast, but very sensitive to artifacts in the EEG. In[1] this is improved by low pass filtering the output of the threshold, but this increases the detection delay (by 200ms), while it is still vulnerable to artifacts. In[2] the line length method from[5] is used, where a detection delay in the order of 4.1 s is reported. Other real-time detection systems that are suitable for closed-loop seizure detection include the work described in[4] based on wavelet filtering combined with median filtering. Depending on the settings, a detection delay of 1.5–2s is achieved. The work reported in[6] features realtime operation, uses an Artificial Neural Network based on Reservoir Computing and has a 1s detection delay. The only FDA approved closed-loop neurostimulation currently available uses several detection mechanisms (bandpass, line-length and area) and reports 600–2000 detections per day[7].

4.3. IMPLEMENTATION

A system overview of the online seizure detection system is depicted in Figure 4.1A. The input of the system consists of the ECoG signals, which are first fed through an analog input filter for signal conditioning that is implemented on a custom PCB. Subsequently the seizure detection filter is realized in the digital domain on a Beaglebone microprocessor. The output can be used to enable seizure suppression techniques, such as for example an optogenetic stimulator. This work focuses on the detection system itself and does not include the stimulator yet. In the following sections the implementation of the proposed system is discussed in detail.

4.3.1. Input filter

The recorded EEG signals are first fed through an analog filter that has four functions: (1) amplification, (2) offset injection in order to match the signal to the input range of the Analog to Digital Converter (ADC), (3) artifact removal by using a second-order 0.4Hz high-pass filter and (4) antialiasing by means of a second-order 23.4Hz low-pass filter. The filter topology that is used for this application is depicted in Figure 4.1B. The first stage is a buffer to prevent the EEG signal from being loaded. Subsequently, the amplitude of the EEG signal is regulated towards $V_{peak\text{-}peak} < 1.8V$, which corresponds to the input range of the ADC. The inverting amplifier accommodates both amplification as well as attenuation in order to handle a large range of input signals. The signal is then passed through a high pass filter to remove unwanted offset and artifacts in the EEG signal. By choosing C_1=2.2 μF and

R_1 = 180 kΩ the cut-off frequency is $f_{c,hpf} = \left(2\pi\sqrt{R_2 R_2 C_2 C_2}\right)^{-1} = 23.4$ Hz. All operational amplifiers are LF353 and use a +5V/-5V power supply.

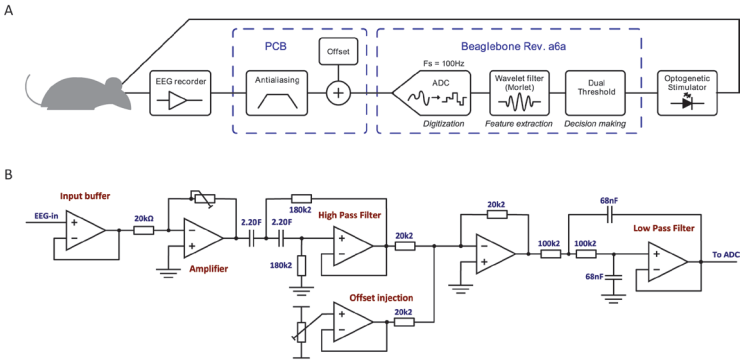

Fig. 4.1 ***(A)*** *the closed-loop epilepsy detection system is depicted with the analog input filter (PCB) and the digital filter (Beaglebone) highlighted.* ***(B)*** *the implementation of the analog input filter is depicted that is used to make the recorded input signal suitable for digitization.*

4.3.2. Digital Wavelet filter

The wavelet filter is implemented on the commercially available Beaglebone microprocessor platform. The processor is an AM335x ARM Cortex-A8 running at 720MHz. The signal is first converted to the digital domain using the onboard 12bit ADC from the Beaglebone using a sampling frequency f_s = 100Hz. The signal is subsequently filtered using a 64th order FIR filter approximating a complex Morlet wavelet filter for which the details are given in the subsequent section. Finally, using an adaptive dual threshold mechanism, a decision is made whether there is a seizure or not.

4.3.3. Complex Morlet for epilepsy detection

In this work the complex Morlet wavelet (Gabor wavelet) was chosen, because of the similarity of the wavelet compared to the morphology of the spikes during a seizure. The complex Morlet is defined as follows:

$$\Psi_\sigma(t) = c_\sigma \pi^{-\frac{1}{4}} e^{-\frac{1}{2}t^2} \left(e^{i\sigma t} - k_\sigma\right)$$

$$k_\sigma = \exp\left(-\frac{1}{2}\sigma^2\right),\ c_\sigma = \sqrt{1 + \exp(-\sigma^2) - 2\exp\left(-\frac{3}{4}\sigma^2\right)}^{-1}$$

and σ determines the scale of the wavelet. To find the optimal value of σ, the response of an ideal continuous-time Morlet wavelet filter is obtained over various frequency scales f_{scale} for each seizure individually using the pre-recorded database. In Figure 4.2 the averaged wavelet output is shown during a seizure. The maximum

output is obtained at f_{scale} = 7 Hz and therefore this scale was chosen for the implementation of the filter.

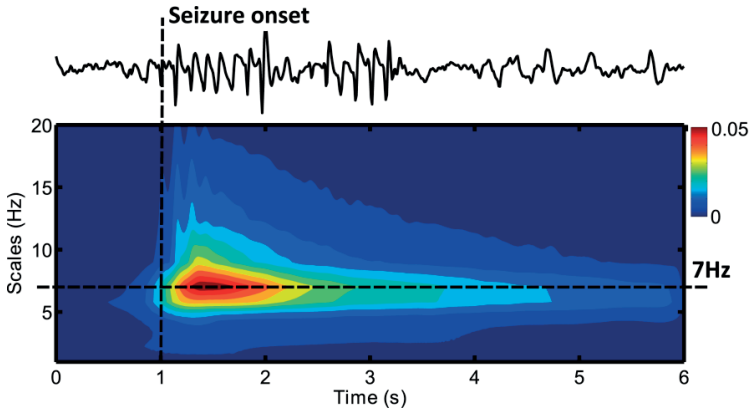

Fig. 4.2 *The average output of a continuous time complex Morlet wavelet filter over 1914 pre-recorded seizures. All seizures have a varying duration, but were aligned to have their onset at 1 s. The maximum output is obtained for the wavelet scale of 7Hz. A single seizure is plotted at the top of the figure for illustrative purposes.*

The ideal wavelet has an infinite time span and therefore has to be truncated when implemented in a practical FIR filter[8]. A trade-off between precision and delay has to be made: the more the wavelet is truncated at the start of the FIR, the faster a detection is made, but the less precise the impulse response approaches the wavelet. To find the optimal truncation, all 1914 seizures were filtered using a set of FIR filters with different truncations. When making a truncation, it was assured that the DC value of the impulse response of the real and imaginary part were both zero by adding an offset, as required for a wavelet filter.

For each truncation setting the average response of the FIR filter was obtained over the dataset. Based on this response the average delay and precision can be determined. The delay of the response was determined by the time at which the maximum in the filter response was reached with respect to the onset of the seizure. This is depicted with the green line in Figure 4.3A. The accuracy is determined by considering how well the filter is able to distinguish between ictal and interictal activity. Therefore the maximum in the averaged response during a seizure is divided by the average value of the response during interictal activity. This is depicted using the black line in Figure 4.3A. For each truncation the ℓ^2 norm of the FIR filter is kept constant: $\sqrt{\sum |x_i|^2} = 1$, in which x_i are the complex coefficients of the wavelet filter.

The accuracy of the filter worsens for truncations > 100, because this is where the impulse response reaches its maximum value. Based on the results in Figure 4.3A, it was chosen to truncate the filter with 101 samples. The accuracy is close to the optimal value, while the delay is 490ms. Note that this is value is not equal to the detection delay of the overall filter, because this also depends on the decision-making strategy that is used. In Figure 4.3B a plot is shown of the FIR with the chosen truncation settings.

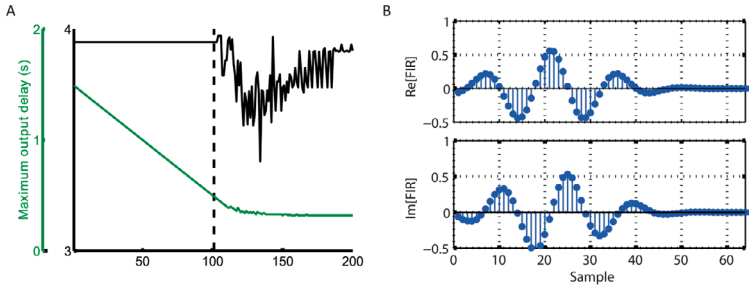

Fig. 4.3 ***(A)*** *the trade-off between detection delay and wavelet precision as a function of the FIR truncation is shown. The green line is the detection delay and the black line shows the accuracy of the wavelet implementation. The black dotted line shows the chosen truncation setting at 101 samples for which the complex (i.e. real and imaginary) impulse response is shown in* ***(B)****.*

Dual Threshold detection

To minimize the detection delay, it was chosen not to include a median filter[4] or other low-pass filters at the output of the wavelet filter. Instead a dual-thresholding mechanism is used: the output is enabled upon a positive crossing of the higher threshold $V_{th,h}$ and is disabled upon negative crossing of the lower threshold $V_{th,h}$ Because of the large variety in signal amplitudes over subjects, the value of both thresholds is determined using an adaptive thresholding technique. $V_{th,l}$ is the running average of the filter output: because the ictal time is short compared to the interictal time, the average will be close to the average of the interictal period. Vth,h is obtained by low-pass filtering the output *Y (i)* of the wavelet filter:

$$V_{th,h}(i) = V_{th,h}(i-1) + \frac{Y(i)-V_{th,h}(i-1)}{\tau fs}$$

$$Y(i) > V_{th,h}(i-1) \rightarrow \tau_{up}$$

$$Y(i) > V_{th,h}(i-1) \rightarrow \tau_{down}$$

This low-pass filtering does not compromise the detection delay, because it is only used to determine $V_{th,h}$ and is not included in the main signal path. By choosing **T**up small, Vth,h will increase rapidly during ictal activity when *Y(i)* > $V_{th,h}$ while a large value of τ_{down} will decrease $V_{th,h}$ slowly during inter-ictal time. In this way

$V_{th,h}$ will reach a value somewhat below the average value during a seizure. In Figure 4.4 the response of the filter during a typical seizure is depicted for illustrative purposes. The effect of the adaptive thresholding is visible, as well as the detection delay, which is around 620ms for this particular seizure. Note that the adaptive thresholding makes the system easy to use compared to systems that need specific training samples[9]. The proposed system is almost 'plug-and-play' since the user only needs to wait for a few seizures before the filter is settled.

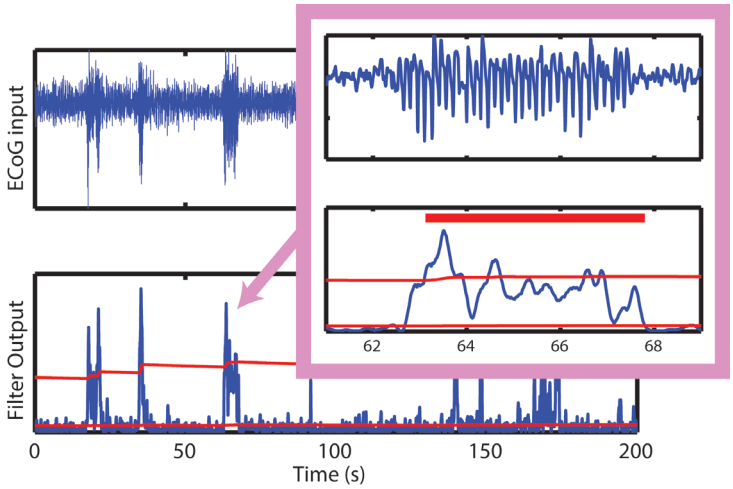

Fig. 4.4 *The input and output of the filter is shown with in the detail the response during the start of a seizure. The red lines correspond to the dual thresholds, while the thick red lines indicate when the output of the filter is enabled.*

4.3.4. Evaluation

The performance of the proposed filter is analyzed over the complete dataset using various values of τ_{up} and τ_{down}. The raw ECoG data was first processed using a filter with a response equivalent to the circuit in Figure 4.1B. The signal is subsequently sampled with f_s = 100Hz and quantized (12 bits accuracy), corresponding with the ADC. Finally, the output of the filter was obtained by replicating the FIR filter and the thresholding mechanism. To allow the values of $V_{th,l}$ and $V_{th,h}$ to settle, the data before the first 4 seizures for each subject from the set of annotated seizures was used as the training dataset and therefore discarded.

A False Negative (FN) was defined as the absence of a detection during an ictal time-stamp from the from the set of annotated seizures. A False Positive (FP) was defined as a detection in the absence of one or more GSWDs in the ECoG. GSWDs were identified using visual inspection of the ECoG by a neuroscientist. This definition was chosen, since upon detection of an GSWD it is not yet known whether this

ictal activity will develop into a seizure or not. In Table 4.1 the performance metrics that are used for comparison are summarized.

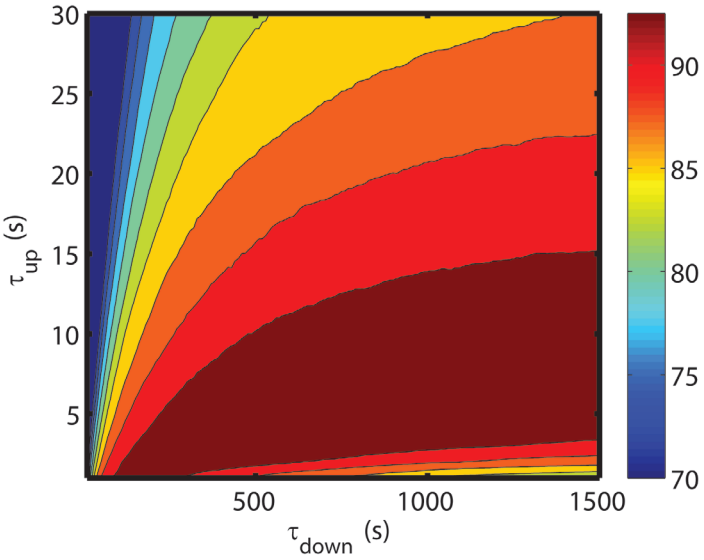

Fig. 4.5 *The Average Detection Rate (ADR) of the filter is given for various values of τ_{up} and τ_{down}.*

Figure 4.5 shows the Average Detection Rate (ADR) of the filter over the complete dataset for various values of τ_{up} and τ_{down}. From this figure the optimal values for a sensitivity of at least 95% were chosen, yielding τ_{up}=5 s and τ_{down}=720 s. The performance of the filter over the complete dataset is summarized in Table 4.2. As can be seen, compared to other implementations, the proposed implementation achieves a very low detection delay at the cost of a decrease in specificity. This trade-off between detection delay and accuracy is very typical for epilepsy detection systems, as was shown in[4]. In the proposed implementation, the feature extraction uses only one component (complex Morlet wavelet) and the decision-making is relatively simple (adaptive thresholding), allowing for a low detection delay and a simple realization with cheap and readily available hardware. Both the feature-extraction as well as the decision-making components of the detection system can benefit from more sophisticated designs, but this will increase the computational complexity and/or the detection delay. To validate the implementation of the filter at the hardware level, the complete system has been realized according to the scheme in Figure 4.1A. It was programmed to generate a 30ms binary pulse upon detection of a seizure, which could be used to trigger a stimulation unit if the system would be deployed in an actual closed-loop setting. After a pulse, the system

will wait 500ms before generating another pulse if the seizure is still active. In Figure 4.6A an example of the output during a seizure from *in vivo* measurement results is shown. In Figure 4.6B it can be seen that the detection delay after onset of this particular seizure is in the order of 200ms. These measurements show that the proposed filter can be used successfully in *in vivo* measurement setups. More measurement data is required to compare the performance of the hardware realization with the results from Table 4.2.

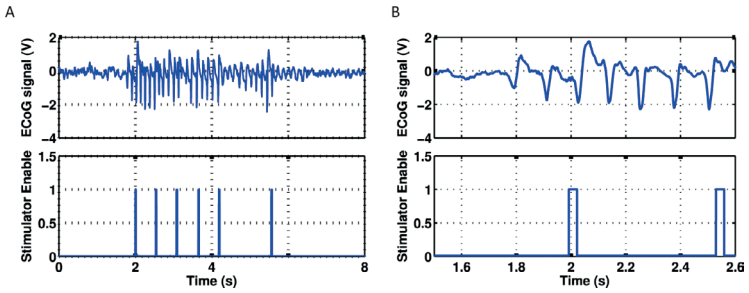

Fig. 4.6 *In vivo measurement results of the epilepsy detection system. After 4 GSWDs in the ECoG signal the detection is made and the stimulator is enabled, which subsequently suppresses the ictal activity.* ***(A)*** *the output during a full seizure is depicted, while in* ***(B)*** *the onset of the seizure is shown.*

4.4. CONCLUSIONS

In this paper the design and realization of an epileptic seizure detection system based on wavelet filtering is described, featuring a very short detection delay. It uses adaptive dual thresholding to accommodate the large variability of input signal levels over multiple subjects without the need to manually change parameters. The system is implemented using off-theshelf components, making it a cheap, flexible and easily reproducible. In vivo measurements confirm the correct functionality in a physical measurement setup.

4.5. REFERENCES

1. Berényi, A., Belluscio, M., Mao, D. & Buzsàki, G. Closed-loop control of epilepsy by transcranial electrical stimulation. *Science*. 337, 735–737 (2012).
2. Paz, J.T., et al. Closed-loop optogenetic control of thalamus as a tool for interrupting seizures after cortical injury. *Nat. Neurosci*. 16, 64–70 (2013).

3. Fanselow, E.E., Reid, A.P. & Nicolelis, M.A., Reduction of pentylenetetrazole-induced seizure activity in awake rats by seizure triggered trigeminal nerve stimulation. *J. Neurosci*. 20, 8160–8168 (2000).
4. Osorio, I., Frei, M.G. & Wilkinson, S.B Real-time automated detection and quantitative analysis of seizures and short-term prediction of clinical onset. *Epilepsia*, 39, 615–627 (1998).
5. Esteller, J., et al. Line length: An efficient feature for seizure onset detection. *Proc. of the 23rd Annual EMBS International Conference*. 1707–1710 (2001).
6. Buteneers, P., et al. Real-time detection of epileptic seizures in animal models using reservoir computing. *Epilepsy Research*. 103, 124–134 (2013).
7. Sun, F.T. & Morrell, M.J. Closed-loop neurostimulation: The clinical experience. *Neurotherapeutics*. 11, 553–563 (2014).
8. Karel, J., et al. Implementing wavelets in continuous-time analog circuits with dynamic range optimization. *IEEE Trans. on Circuits and Systems – I*. 59, 229–242 (2012).

4.6. TABLES

Table 4.1 Performance metrics used to benchmark the proposed detection system

Metric	Description
Detection delay	The delay between the onset of the seizure annotation and the detection of the filter.
FPPS	False-Positive-Per-Seizure rate: the number of False positives divided by the number of seizures
FNPS	False-Negative-Per-Seizure rate: the number of False negatives divided by the number of seizures.
Sensitivity	The percentage of seizures that was detected successfully
Specificity	The percentage of correctly classified inter-ictal intervals.
ADR	Average Detection Rate: 0.5(sensitivity + specificity)

Table 4.2 Comparison of the performance of various real-time seizure detection systems

Reference Metric	5	4	6	This work
Detection delay ± SD	4.1±3.8 s	1.5±7.2 s	0.97±0.93 s	0.49±0.16 s
FPPS	1.44	0.04	0.09	0.09
FNPS	0.02	0	0.07	0.04
Specificity	Unknown	Unknown	98.2%	93.6%
Sensitivity	Unknown	100%	96.2%	96.0%
ADR	Unknown	Unknown	91.2%	94.8%