

Implementing Wavelets in Continuous Time Analog Circuits with Dynamic Range Optimization

Citation for published version (APA):

Karel, J. M. H., Haddad, S. A. P., Hiseni, S., Westra, R. L., Serdijn, W. A., & Peeters, R. L. M. (2012). Implementing Wavelets in Continuous Time Analog Circuits with Dynamic Range Optimization. *IEEE Transactions on Circuits and Systems*, 59(2), 229-242. <https://doi.org/10.1109/TCSI.2011.2162381>

Document status and date:

Published: 01/02/2012

DOI:

[10.1109/TCSI.2011.2162381](https://doi.org/10.1109/TCSI.2011.2162381)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Implementing Wavelets in Continuous-Time Analog Circuits With Dynamic Range Optimization

Joël M. H. Karel, *Member, IEEE*, Sandro A. P. Haddad, *Member, IEEE*, Senad Hiseni, *Student Member, IEEE*, Ronald L. Westra, Wouter A. Serdijn, *Fellow, IEEE*, and Ralf L. M. Peeters

Abstract—Signal processing by means of analog circuits offers advantages from a power consumption viewpoint. A method is described to implement wavelets in analog circuits by fitting the impulse response of a linear system to the time-reversed wavelet function. The fitting is performed using local search involving an L_2 criterion, starting from a deterministic starting point. This approach offers a large performance increase over previous Padé-based approaches and allows for the circuit implementation of a larger range of wavelet functions. Subsequently, using state-space optimization the dynamic range of the circuit is optimized. Finally, to illustrate the design procedure, a sixth-order L_2 -approximated orthonormal Gaussian wavelet filter using G_m -C integrators is presented.

Index Terms—Analog circuits, design methodology, dynamic range, filters, low power, signal processing, wavelet filters, wavelets, wavelet transform.

I. INTRODUCTION

THE wavelet transform [1], [2] is a widely used signal processing technique that is becoming more and more common in medical signal processing applications, especially in cardiac signal processing [3]–[9]. The main advantage of this technique over the classical Fourier transform is that it provides combined time and frequency localization, whereas in the Fourier domain only frequency information is available. Also, the wavelet transform is more flexibly employed than the Fourier transform, as it can be tuned to the typical signal morphology and particularities of an application at hand: instead of having to work with harmonic basis functions, as is the case for the Fourier transform, a large class of wavelet functions is available that obey the wavelet admissibility conditions. There are two forms of the wavelet transform: 1) the *discrete-time*

wavelet transform and 2) the *continuous-time* wavelet transform, both employing continuous-time wavelet functions. In this paper, we shall be concerned with the continuous-time wavelet transform, which involves a moving L_2 -inner product between the signal and the scaled wavelet function, usually for a continuous range of multi-resolution scales (continuous dilation). For an in-depth discussion on wavelet transforms the reader is referred to [2].

In wearable and implantable medical devices, such as pacemakers, the wavelet transform potentially has a large number of useful applications [9]. However, such devices impose strict constraints on the power consumption [10]–[17]. Especially for the sensing circuit of pacemakers, which is always active and for which the battery can only be replaced after surgical intervention, power consumption is an important factor. When implementing the discrete-time wavelet transform in a digital fashion, analog-to-digital (A/D) converters are required. These A/D converters account for a relatively high power consumption, that increases exponentially with the number of bits used by the A/D converters. From a power consumption perspective, an obvious alternative is to use analog circuits and to implement the continuous-time wavelet transform instead.

The concept of implementing a single scale of the continuous-time wavelet transform by constructing a linear system, the impulse response of which matches a time-reversed and time-shifted Gaussian wavelet function, was introduced in [18] and the term “wavelet filter” was coined for the first time in [10]. Over the years, this idea has been worked out in more detail, concentrating on issues related to wavelet approximation [11], [19]–[21], wavelet filter state-space and wavelet filter topology optimization [21]–[24] and wavelet filter circuit design [10], [22], [25]–[27]. In this paper, for the first time, a comprehensive treatment of how to map any continuous-time wavelet transform on a low-order wavelet filter topology that is optimized for dynamic range, sparsity and sensitivity is presented.

This paper is structured as follows. In Section II, it is argued how a single scale of the continuous-time wavelet transform can be implemented with a linear system. In Section III, a novel L_2 -approximation approach is presented, which allows one to obtain an accurate low order system approximation for many different wavelets, to be used for hardware implementation. Section IV addresses circuit design issues, such as optimization of the dynamic range, the tradeoff between a high dynamic range and sparsity, and sensitivity. Finally, to illustrate the design procedure, a sixth-order L_2 -approximated orthonormal Gaussian (gaus1) wavelet filter using G_m -C integrators is presented in Section V.

Manuscript received October 06, 2009; revised January 07, 2011, April 13, 2011 and June 14, 2011; accepted June 20, 2011. Date of publication August 30, 2011; date of current version January 27, 2012. This work is part of the BioSens project supported by STW, the Dutch Technology Foundation. This paper was recommended by Associate Editor A. J. Lopez-Martin.

J. M. H. Karel, R. L. Westra, and R. L. M. Peeters are with the Department of Knowledge Engineering, Maastricht University, 6200 MD Maastricht, The Netherlands (e-mail: joel.karel@maastrichtuniversity.nl).

S. A. P. Haddad is with the Faculty of Gama (FGA), University of Brasília, 71405-610 Gama-DF, Brazil (e-mail: sandrohaddad@unb.br).

S. Hiseni and W. A. Serdijn are with the Electronics Research Laboratory, Delft University of Technology, Delft 2628 CD, The Netherlands (e-mail: w.a.serdijn@tudelft.nl).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCSI.2011.2162381

II. WAVELETS IN ANALOG CIRCUITS

At a given time $\tau \in \mathbb{R}$ and scale $\sigma \in \mathbb{R}^+$, and for a given admissible wavelet function $\psi(t)$, the value of the continuous-time wavelet transform of a real signal $f(t)$ is defined as the cross-covariance at lag τ of that signal and the normalized dilated wavelet $(1/\sqrt{\sigma})\psi(t/\sigma)$ [2]:

$$W_\psi(\tau, \sigma) = \frac{1}{\sqrt{\sigma}} \int_{-\infty}^{\infty} f(t) \psi\left(\frac{t-\tau}{\sigma}\right) dt. \quad (1)$$

Typically, an admissible wavelet function $\psi(t)$ is required to (virtually) have compact support: it effectively vanishes outside a certain bounded interval. Then the time shift τ becomes naturally associated with time localization. Also, the spectrum of $\psi(t)$ typically is bounded, exhibiting a dominant frequency, say f_c , called the center frequency of the wavelet. Then at the scale σ , the center frequency of $\psi((t-\tau)/\sigma)$ is given by σf_c ; thus, the scale σ is naturally associated with this pseudo-frequency. The scales σ allow for zooming in on interesting features of the signal associated with typical frequencies, while maintaining time localization to a certain extent. For practical implementation in analog circuits, the wavelet transform is usually computed for continuous values of τ but only at a few scales σ of interest.

When a signal $f(t)$ is passed through a linear time-invariant (LTI) filter, the filter output is well known to be the convolution of $f(t)$ with the impulse response $h(t)$ of the filter:

$$(f * h)(\tau) = \int_{-\infty}^{\infty} f(t) h(\tau - t) dt. \quad (2)$$

From (1) and (2) it is immediate that analog computation of the wavelet transform $W_\psi(\tau, \sigma)$ at a selected (fixed) scale σ , is achievable through implementation of a linear filter with impulse response

$$h(t) = \frac{1}{\sqrt{\sigma}} \psi\left(\frac{-t}{\sigma}\right). \quad (3)$$

Equivalently, the transfer function of the filter needs to satisfy

$$H(s) = \frac{1}{\sqrt{\sigma}} \int_{-\infty}^{\infty} \psi\left(\frac{-t}{\sigma}\right) e^{-st} dt. \quad (4)$$

For a given wavelet function $\psi(t)$ this ideal transfer function $H(s)$ will usually be non-rational (so that the associated filter is infinite dimensional) and non-causal (requiring future values of the signal $f(t)$ to compute the value of the wavelet transform at a given point in time) but stable (due to the wavelet admissibility condition). For an application in practice; however, a reasonably good *approximation* of $H(s)$ may be sufficient to achieve the intended goals.

For obvious physical reasons we shall only consider the hardware implementation of *finite dimensional, strictly causal, stable LTI filters* of sufficiently low order. Thus, we consider approximations of $H(s)$ by transfer functions that have all their poles in the left half of the complex plane and which are strictly




Fig. 1. Effect of time-shifting.

proper rational (i.e., the degree of the numerator polynomial is strictly less than the degree of the denominator polynomial).

Causality of a filter implies that $h(t)$ will be zero for negative t , so that any time-reversed mother wavelet $\psi(-t)$ which does not have this property must be *time-shifted* by some suitable value t_0 , to facilitate an accurate approximation of its correspondingly time-delayed wavelet transform [22] $\tilde{W}_\psi(\tau, \sigma)$

$$\tilde{W}_\psi(\tau, \sigma) = \frac{1}{\sqrt{\sigma}} \int_{-\infty}^{\infty} f(t) \tilde{\psi}\left(\frac{\tau-t}{\sigma}\right) dt \quad (5)$$

where

$$\tilde{\psi}(t) = \psi(t_0 - t). \quad (6)$$

In case $\tilde{\psi}(t)$ is nonzero on the negative t -axis, a truncation error results, which should be kept small. This is illustrated in Fig. 1. The selection of the time-shift t_0 turns out to be a rather delicate process; see [11]. When t_0 is chosen too small, the truncation error may be too large, exerting a significant decrease of the overall performance of the approximation. When t_0 is chosen too large, then an accurate $h(t)$ needs to remain close to zero for some time, requiring the implementation of a high-order filter. This is undesirable, as the complexity of the design increases with the order of the filter, and so do the physical space required on the chip and the associated power consumption.

III. L_2 -APPROXIMATION OF WAVELET FUNCTIONS

A. Motivation of an L_2 -Approximation Criterion

For a selected time-shift t_0 , the computation of an accurate approximation is a nontrivial task in itself, which allows for various approaches with varying success. In earlier work, Padé approximation (see [28], [29]) was used to approximate the Laplace transform of $\psi(t)$ with a strictly proper rational function $H(s)$; see [22]. While a Padé approximation is easy to compute (uniquely) by solving a linear system of equations, a main disadvantage of this approach is that many well-known wavelets do not have a known expression for their Laplace transform. Also, the choice of degrees for a sufficiently accurate Padé approximation is not straightforward and hard to automate. Only a very limited set of wavelets, including the

Gaussian wavelet and the Mexican Hat wavelet, could be handled satisfactorily in this way. Note that the Padé approach with a good fit in the Laplace domain near zero, corresponds to a good fit in the time domain at infinity (final value theorem) and hardly constrains the wavelet on its domain of support, often yielding a poor fit. Requiring a good fit in the Laplace domain at infinity yields a much better fit for the wavelet on a part of its support near zero (initial value theorem) but often results in instability. This tradeoff is not handled well by the Padé approach [11]. In [21], [24], the Padé approximation was implemented using the so-called CFOS-approach: by using a cascade of complex first order systems (CFOS), see [30], [31], one may model a Gaussian envelope of a wavelet function separately from its oscillatory characteristics. Using this CFOS-approach the Gabor wavelet and the Morlet wavelet could successfully be approximated and implemented in hardware. However, this method is again limited to a small set of wavelets and the filter order is not easily kept low.

In [11], preliminary results on a novel, more accurate L_2 -approach have been reported. This method proceeds by fitting the impulse response $h(t)$ of a filter directly to a given function $\tilde{\psi}(t)$ by minimizing the L_2 -norm of the error function $h(t) - \tilde{\psi}(t)$ (for $t \geq 0$). Here we present a detailed account of this L_2 -approach, including a methodology which facilitates largely automated approximation.

To motivate the use of an L_2 -criterion to quantify the error between an approximation $h(t)$ and a given (shifted, truncated) mother wavelet $\tilde{\psi}(t)$, it is noted that: 1) the value of the wavelet transform at a given location and scale involves an L_2 -inner product; 2) to compute the wavelet transform on a given scale, the function $\tilde{\psi}(t)$ is used as a convolution kernel with arbitrary time-shifts, so that approximation errors for all time instances t have an equal impact on the overall accuracy of the results; 3) due to Parseval's theorem, L_2 -approximation in the time domain can be reformulated equivalently as L_2 -approximation in the frequency domain, which allows one to deal with the transfer function of a wavelet filter directly (cf. [11]). In fact, it holds that

$$\left| W_h(\tau, \sigma) - W_{\tilde{\psi}}(\tau, \sigma) \right| \leq \|f(t)\| \|h(t) - \tilde{\psi}(t)\| \quad (7)$$

see [20, Thm. 4.4.1] so that the L_2 -norm of the error $h(t) - \tilde{\psi}(t)$ can be used to establish an *a priori* bound on the absolute accuracy of the wavelet transform.

In practice, it may happen that no analytical expression is available for the function $\tilde{\psi}(t)$, which is the case, e.g., for the Daubechies wavelets and many other commonly used classes of wavelets. Then, by sampling $\tilde{\psi}(t)$ with a small sampling time over an interval which is sufficiently large, the L_2 -criterion of fit may be conveniently approximated by a nonlinear sum of squares, for which numerical optimization software is widely available. We have used this technique to obtain the numerical results presented in this paper.

It is noted that a well-known drawback of the use of nonlinear least squares techniques is the possible existence of local (non-global) optima. An iterative local search optimization routine may terminate in such a local optimum and as a rule there

is no guarantee whether a global optimum will be found, nor will it be possible to ascertain global optimality once a global optimum has actually been found. It has been found that lower order systems sometimes give better results for this very reason, and different local search algorithms provide different local optima [20]. A large amount of different starting points can yield different local optima and thus can be used to find better solutions, at a relatively high computational cost. The choice of a clever initial point [19] may have the same effect, but with considerably lower computational costs. A procedure to find a good initial point for the present application is discussed in Section III-C.

B. Parameterization

To minimize the L_2 -norm of the error function $h(t) - \tilde{\psi}(t)$ numerically, a parameterized class of functions must be specified over which to optimize for $h(t)$. From the theory of linear systems, see [32], it is known that any strictly causal LTI filter of finite order n can be represented as a state-space system (A, B, C) , corresponding to a system of associated first-order differential equations:

$$\frac{dx(t)}{dt} = Ax(t) + Bu(t) \quad (8)$$

$$y(t) = Cx(t). \quad (9)$$

Here, the filter input is denoted by $u(t)$ (in the application at hand this concerns the signal $f(t)$), the output is denoted by $y(t)$ (which represents the wavelet transform of $f(t)$ at a given scale) and $x(t) \in \mathbb{R}^n$ denotes the internal n -dimensional state vector at time t (which may or may not have a direct physical or electrical interpretation, depending on the possible structure of A , B , and C). In this general black-box description, the associated impulse response $h(t)$ and its Laplace transform (the transfer function) $H(s)$ of the system are given by

$$h(t) = Ce^{At}B \quad (10)$$

$$H(s) = C(sI_n - A)^{-1}B. \quad (11)$$

For the generic situation of stable systems with distinct poles, the impulse response $Ce^{At}B$ is a linear combination of damped exponentials and exponentially damped harmonics. For low-order filters, this makes it possible to propose an explicitly parameterized class of impulse response functions among which to search for a good approximation of $\tilde{\psi}(t)$; see also [11], [19]. As an example, a useful class of impulse responses of order $n = 5$ consists of linear combinations of a single damped exponential and two exponentially damped harmonic pairs:

$$h(t) = \alpha_1 e^{p_1 t} + \beta_1 e^{q_1 t} \sin(r_1 t) + \gamma_1 e^{q_1 t} \cos(r_1 t) + \beta_2 e^{q_2 t} \sin(r_2 t) + \gamma_2 e^{q_2 t} \cos(r_2 t). \quad (12)$$

Note that wavelets are typically oscillatory functions which effectively have compact support, so that any good approximation necessarily requires a number of damped harmonics. To have stability the parameters p_1 , q_1 and q_2 must all be strictly negative.

For the purpose of (integrated) circuit design, it is useful to have an explicit state-space representation (A, B, C) too. For $h(t)$ of the form (12) such a representation is provided by

$$\begin{aligned} A &= \begin{pmatrix} p_1 & 0 & 0 & 0 & 0 \\ 0 & q_1 & r_1 & 0 & 0 \\ 0 & -r_1 & q_1 & 0 & 0 \\ 0 & 0 & 0 & q_2 & r_2 \\ 0 & 0 & 0 & -r_2 & q_2 \end{pmatrix} & B &= \begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \\ 1 \end{pmatrix} \\ C &= (\alpha_1 \quad \beta_1 \quad \gamma_1 \quad \beta_2 \quad \gamma_2). \end{aligned} \quad (13)$$

For other values of the order n , similar parameterizations are easily generated in an analogous fashion.

Given the explicit form of the wavelet $\tilde{\psi}(t)$ and the parameterized class of functions $h(t)$, the (squared) L_2 -norm of the error function $h(t) - \tilde{\psi}(t)$ can now be minimized in a straightforward way using standard numerical optimization techniques and software. When required (e.g., when the wavelet is only available in sampled numerical form), one may employ a discretized problem setting obtained by sampling both the wavelet and the impulse response with a suitably small sampling time. The negativity constraints on p_1 , q_1 and q_2 which enforce stability are easily handled by most optimization packages.

One important admissibility property of wavelets that was undiscussed so far, is that it needs to have at least one *vanishing moment*. This comes down to the requirement that the integral of the wavelet function equals zero:

$$\int_{-\infty}^{\infty} \tilde{\psi}(t) dt = 0. \quad (14)$$

If the approximation $h(t)$ to $\tilde{\psi}(t)$ does not share this property, this will generally cause an unwanted time-varying bias in the approximated wavelet transform of $f(t)$. In fact, if no special care is taken, this is likely to happen in a situation where a truncation error occurs. Therefore, the additional constraint needs to be imposed for the approximation $h(t)$ to have a zero integral. In terms of the transfer function $H(s)$, this condition is equivalently captured by the constraint

$$H(0) = -CA^{-1}B = 0 \quad (15)$$

which expresses that the filter needs to have a zero at $s = 0$. For the explicit parameterization introduced above, in which A is block diagonal, the inverse A^{-1} can easily be determined. As a concrete example, for the parameterization (12), the following explicit condition results:

$$\frac{\alpha_1}{p_1} + \frac{-\beta_1 r_1 + \gamma_1 q_1}{q_1^2 + r_1^2} + \frac{-\beta_2 r_2 + \gamma_2 q_2}{q_2^2 + r_2^2} = 0. \quad (16)$$

This condition needs to be passed on to the optimization software. If such a nonlinear constraint is not conveniently handled by the optimization software, one may instead choose to eliminate a parameter to enforce it [20].

C. Finding a Good Starting Point

The possible existence of local optima for the L_2 -approximation criterion, makes that the choice of a starting point may

have a considerable impact on the outcome of a local search optimization routine. In [19], an approach was introduced which aims to find a good starting point quickly. This method employs standard model reduction techniques to compute a low-order approximation from an accurate initial high order approximation. The resulting overall L_2 -approximation method can then largely be automated. It consists of a number of steps, as illustrated in Fig. 2.

1) Sampling the wavelet $\tilde{\psi}(t)$

The function $\tilde{\psi}(t)$ is sampled with high resolution over a sufficiently large time interval, yielding a discrete-time signal $\{g_k\}$

$$g_k = \tilde{\psi}(k\Delta t), \quad k = 0, 1, \dots, N-1 \quad (17)$$

where Δt denotes the sampling time and N is the total number of sampling points. This sampled sequence acts as a highly accurate discrete-time representation of the waveform that needs to be implemented as the impulse response $h(t)$ of a linear filter. Choosing the time interval to be large enough effectively promotes stable approximations to occur in further computations, in view of the fact that a wavelet quickly vanishes when t becomes large.

2) Computation of a low-order approximation, used as a starting point for optimization

This involves four steps, each of which are supported by widely available software packages, such as the Control System Toolbox of Matlab.

First, a *high order discrete-time FIR model* is constructed of which the impulse response $\{h[k]\}$ coincides with the sampled sequence. More precisely, it is required that $h[0] = 0, h[1] = g_0, \dots, h[N] = g_{N-1}$ and in addition $h[k] = 0$ for all $k > N$. A state-space realization in controllable companion form is convenient here. Typically, in our set-up, the order of this high-order FIR model equals about $N = 2000$, with $\Delta t = 0.01$.

Second, the high order state-space model is *balanced and truncated*, see [33] and [34], to yield a still highly accurate reduced order discrete-time model. Typically, the order of this “intermediate order model” is between 50 and 100.

Third, this discrete-time intermediate order model is converted to a *continuous-time* model, e.g., using the zero-order hold (ZOH) principle.

Finally, the intermediate order continuous-time model is reduced with a balance and truncate technique to a desired low order n , typically between 3 and 15. If desired, various low orders n can be tried out and compared and a suitable low order of a satisfactory approximation quality can be selected. This is fast, as all the previous steps need not be repeated.

3) L_2 -approximation of the wavelet function

The low-order model obtained in the previous step is used as a starting point for minimizing the L_2 -norm of the error function $h(t) - \tilde{\psi}(t)$ under the constraint (15), using an iterative local search optimization technique. To obtain the initial parameter values, the low-order approximation model must be cast into the correct state-space form first, which

TABLE I
NORMALIZED L_2 -ERROR OF VARIOUS APPROXIMATIONS, EXCLUDING THE POTENTIAL TRUNCATION ERROR

Wavelet	order 12	order 10	order 8	order 7	order 6	order 5	order 3
Gaussian	0.0044	0.0103	0.0056	0.0070	0.0132	0.0475	0.3976
Morlet	0.0028	0.0040	0.0294	0.0676	0.0984	0.2520	0.6712
Mexican Hat	0.3791	1.0551	1.1361	3.0083	0.1108	1.1011	0.5020
Daubechies 3	0.1184	0.5891	0.6484	0.3791	0.3796	0.4540	0.9999
Daubechies 7	0.0369	0.0429	0.0993	0.2430	0.2510	0.5241	0.8400
Coiflet 5	0.0284	0.0706	0.1832	0.3510	0.4340	0.6335	0.8841




Fig. 2. Automated approximation of wavelet functions.




Fig. 3. Approximations of various wavelet functions. The thick gray line corresponds to the wavelet function and the thin dashed black line to the impulse response of the approximating system. The following wavelets were approximated with specified orders $O(\cdot)$ and time shifts $D(\cdot)$. (a) Gaussian $O(5)$, $D(2)$. (b) Morlet $O(5)$, $D(2.5)$. (c) Mexican Hat $O(6)$, $D(3.2)$. (d) Daubechies 3 $O(6)$, $D(-1)$. (e) Daubechies 7 $O(8)$, $D(-4)$. (f) Coiflet 5 $O(8)$, $D(-11)$.

is not difficult to achieve. Note that a numerical optimization approach to the minimization problem $\min \int_0^\infty (h(t) - \hat{\psi}(t))^2 dt$ involves a discretization too, which can be carried out with a resolution that is different from the resolution used for sampling in Step 1.

In the procedure for computing a low-order starting point, we start from an accurate high-order approximation (the FIR model) which is first reduced in discrete-time, then converted to continuous-time, and finally reduced in continuous-time to a desired low order. One cannot convert the initial FIR model directly to continuous-time, as this involves a logarithmic operation and all the system poles are—by construction—located at the origin. These poles are discarded and moved away from the origin in the discrete-time model reduction step, making the subsequent steps feasible. Of course, any other approach to

compute a good initial approximation can be used to replace this procedure in the overall scheme.

This L_2 -approach allows for the approximation and implementation of a wide variety of wavelets. In [24], it was discussed how certain complex wavelets can be implemented and in [20] various other wavelets are approximated with this approach. Besides the Gaussian wavelet, also the Morlet wavelet and the Mexican Hat wavelets have been approximated, all of which are of interest for ECG processing; see [35]–[38]. The current approach also allows for the approximation of wavelets that have an associated filter bank such as the Daubechies and Coiflet wavelets. The approximated wavelets are shown in Fig. 3.

In [19] and [20], it is shown that an L_2 -approximation of order 4 of a Gaussian wavelet outperforms a Padé approximation of order 5. The adopted L_2 -approximation methodology is capable of yielding a significant improvement in approximation performance together with a further reduction of the approximation order. To illustrate how well various wavelets are approximated, the L_2 -norms of the approximation errors of various wavelets are displayed in Table I. Note that a prescribed set of orders was used for all the wavelet approximations.

Accuracy does not always increase with the order. One of the reasons is that the configuration, i.e., the ratio between the number of exponentially damped harmonics and damped exponentials, may be important, especially when only one deterministic starting point is used. If a local optimum happens to be nearby, then the optimization may quickly terminate there; this occurred for the Mexican Hat wavelet, for several orders. Switching to another local search technique such as a trust-region reflective search solves this particular problem; see [20], but then the results for the Daubechies wavelets deteriorate. Another way to overcome this problem, at the cost of a longer computation time, is to use multiple random starting points in an area around the deterministic one.

The L_2 -approach to wavelet approximation was compared with the Padé and CFOS approaches in terms of the mean-square error (MSE); see Fig. 4. Since not all approaches can handle all wavelets, the MSE was calculated for the Gaussian wavelet and for the Morlet wavelet, using several filter orders. The L_2 -approach typically achieves a specified accuracy with a lower order system than the Padé approach; the CFOS requires a substantial amount of stages to achieve an accurate approximation to the Gaussian envelope. For a limited complexity (order of the filter) the L_2 -approach provides the best accuracy.

IV. STATE-SPACE OPTIMIZATION

This section describes the next step towards the implementation of the wavelet filter, i.e., the definition of its wavelet




Fig. 4. Order of the filter versus mean-square error for L_2 , Padé and CFOS approximations. (a) First derivative of Gaussian and (b) Morlet wavelet bases.

filter topology. Since state-space descriptions and their corresponding filter topologies are not unique representations of a dynamical system, they allow the designer to find an implementation that fits the requirements best, e.g., easy coefficients, a prescribed filter topology, or maximum dynamic range. The state-space description is transformed into the desired form by state-space transforms or similarity transforms. In the context of low-power, low-voltage analog integrated circuits, the most important requirements are dynamic range, sensitivity and sparsity, all of which will be treated in the subsections that follow.

A. Dynamic Range Optimization

A system's dynamic range is essentially determined by the maximum processable signal magnitude and the internally generated noise. It is well known that the system's controllability and observability gramians play a key role in the determination and optimization of the dynamic range [39], [40].

The controllability and observability gramians are derived from the state-space description. If A is stable, the controlla-

bility (K) and observability (W) gramians are the unique, symmetric solutions of the following two Lyapunov equations:

$$AK + KA^T + 2\pi BB^T = 0 \quad (18)$$

$$A^T W + WA + 2\pi C^T C = 0 \quad (19)$$

where A , B and C are the state, input, and output matrices of the state-space description, respectively.

In [23], it is shown that, in order to maximize the dynamic range of the system, one should minimize the objective functional F_{DR} that represents the relative improvement of the dynamic range and contains all parameters that are subject to manipulation by the designer. The objective functional is given by

$$F_{DR} = \frac{\max_i k_{ii}}{(2\pi)^2} \sum_i \frac{\alpha_i}{C_i} w_{ii} \quad (20)$$

where k_{ii} and w_{ii} are the main diagonal elements of K and W , respectively, $\alpha_i = \sum_j |A_{ij}|$ is the absolute sum of the elements on the i th row of A , and C_i is the capacitance in integrator i .

The optimization of F_{DR} employs state-space transforms, also known as similarity transforms. A similarity (coordinate) transform, involving an invertible transformation matrix T , defines a new state vector

$$x' = T^{-1}x. \quad (21)$$

Consequently, the transformation results in new system matrices, given by

$$\begin{aligned} A' &= T^{-1}AT \\ B' &= T^{-1}B \\ C' &= CT \end{aligned} \quad (22)$$

Using the above state-space transformation, the controllability and observability gramians can be optimized as they become

$$\begin{aligned} K' &= T^{-1}KT^{-T} \\ W' &= T^T WT \end{aligned} \quad (23)$$

where $T^{-T} = (T^{-1})^T$.

As the dynamic range of a circuit is defined as the ratio of the maximum and the minimum signal level that it can process, optimization of the dynamic range is equivalent to the simultaneous maximization of the (distortionless) output swing and the minimization of the overall noise contribution [41].

Thus, the first optimization step boils down to finding a similarity transform, such that the controllability gramian of the new system becomes a diagonal matrix with equal diagonal entries. Such a transform is given by

$$T_K = P_K D_K^{1/2} \quad (24)$$

where P_K is the eigenvector matrix of K and D_K is a diagonal matrix, the diagonal entries of which are the eigenvalues of K .

In the second step of the optimization procedure, the system is optimized with respect to its noise contribution. While preserving the result of the first optimization step, it is possible to

rotate the state space, such that the observability gramian becomes a diagonal matrix as well. The transformation of W' , leading to W'' , is defined as

$$T_W = P_{W'} \quad (25)$$

where $P_{W'}$ is the eigenvector matrix of W' .

Finally, profiting from the well-known fact that the relative noise contribution of an integrator decreases when the capacitance and bias current increase, we match an optimal capacitance distribution to the noise contributions of each individual integrator (noise scaling), i.e., the diagonal entries of W'' , according to [23]

$$C_i = \frac{\sqrt{\alpha_i w''_{ii} k''_{ii}}}{\sum_j \sqrt{\alpha_j w''_{jj} k''_{jj}}} \cdot C_{\text{tot}} \quad (26)$$

C_{tot} being the total capacitance of the wavelet filter.

B. Sparsity Versus Dynamic Range Comparison

The drawback of a dynamic-range optimal system is that its state-space matrices are generally fully dense, i.e., all the entries of the A , B , and C matrices are filled with nonzero elements. These coefficients will have to be mapped on circuit components, and will result in a complex circuit with a large number of interconnections. For high-order filters it is therefore necessary to investigate how a realization of the desired transfer function having sparser state-space matrices would compare to the one having maximal dynamic range.

By definition, a sparse matrix is a matrix populated primarily with zeros. For a less complex circuit, it is possible, for instance, to transform the state-space matrix A by a similarity transformation matrix T in order to reduce the number of non-zero coefficients in A . One important characteristic of T is that, considering a system of order n , T has n^2 degrees of freedom. Thus, the system can be transformed in several ways, achieving the desired sparsity requirements.

In order to compare the dynamic range versus the order of the system, Fig. 5 illustrates the F_{DR} of different state-space descriptions applied to a Morlet wavelet filter transfer function. Note that for lower order, all state-space representations present a similar performance, with F_{DR} values close to each other, whereas for higher order, the F_{DR} 's of some representations deviate too much from the optimal case. Hence, dynamic range optimization becomes a very important requirement in high-order analog wavelet filter design. As one can see, the orthonormal ladder state-space representation [42] is very close to the optimal representation, even for high-order systems.

C. Sparsity Figure-of-Merit (SFOM)

In order to define the sparsity property versus dynamic range for a specific state-space description, we derive a new figure-of-merit. The static power consumption of an analog active filter is basically determined by the bias currents inside the filter. From a state space representation one can define the total bias current by $N_c I_o$ where N_c represents the number of nonzero coefficients presented in matrices A , B and C and I_o is the current




Fig. 5. F_{DR} versus order of a Morlet wavelet filter for different state-space representations.

necessary to implement each coefficient.¹ As a result, the static power consumption of a filter can be approximated as

$$P_{\text{static}} = N_c I_o V_{cc} \quad (27)$$

On the other hand, the dynamic power dissipation of a n th order filter can be expressed in terms of its dynamic range, DR, and yields [23]

$$P_{\text{dynamic}} = 8n f k T \xi \text{DR} = 4n f \left(\frac{M}{\delta(p)} \right)^2 \text{Tr}(KQ) \frac{1}{F_{\text{DR}}} \quad (28)$$

where n is the order of the filter, f is the operating frequency, k is Boltzmann's constant, T is the absolute temperature, ξ is the noise factor of the integrator, which is fundamentally greater or equal than $1/2$, M is the maximum output amplitude, and $\delta(p)$ is a nonlinear monotonically increasing function of the fraction of time p that the integrators are allowed to clip and is represented by $\delta(p) = -2 \ln(1 - p)$. Tr represents the trace of a matrix and $Q = C^T C$ is the state weighing matrix. $\text{Tr}(KQ)$ is invariant under a similarity transformation and is thus beyond the designer's control.

Dividing P_{static} by P_{dynamic} we end up with

$$\begin{aligned} \frac{P_{\text{static}}}{P_{\text{dynamic}}} &= \frac{N_c I_o V_{cc}}{4n \left(\frac{M}{\delta(p)} \right)^2 \text{Tr}(KQ) \frac{1}{F_{\text{DR}}}} \\ &= \frac{I_o V_{cc}}{4 \left(\frac{M}{\delta(p)} \right)^2 \text{Tr}(KQ)} \frac{N_c F_{\text{DR}}}{n} \end{aligned} \quad (29)$$

where $(I_o V_{cc}) / (4(M/\delta(p))^2 \text{Tr}(KQ))$ is defined by the circuit implementation, while $N_c F_{\text{DR}} / n$ can be controlled at the system level. So, in order to design an optimal state-space description with respect to dynamic range and sparsity, we present a new figure-of-merit, given by

$$\text{SFOM} = \frac{N_c \times F_{\text{DR}}}{n} \quad (30)$$

¹For simplicity, it is assumed that every coefficient to be implemented entails a current consumption of I_o .

TABLE II
NUMBER OF NONZERO COEFFICIENTS N_c IN A , B , AND C

Optimal	$n^2 + 2n$
Hessenberg	$\frac{n(n+1)}{2} + 3n - 1$
Schur	$\frac{n(n+1)}{2} + 2n$
Canonical	$3n$
Biquad	$\lfloor \frac{1}{2}n^2 + 3n \rfloor$
Orthonormal	$3n$




Fig. 6. SFOM versus order of the Morlet wavelet filter for different state-space representations.

From (30) one can see that the optimal case is the one with minimum SFOM, i.e., the system which presents a good sparsity (large number of zeros) while preserving a low F_{DR} (i.e., a large dynamic range).

The number of nonzero coefficients N_c for different state-space representations is given in Table II.

Finally, several SFOMs versus the order of the system are plotted in Fig. 6. As expected, the orthonormal ladder state-space representation presents the best performance (lower SFOM), mainly due to its excellent sparsity and its near-optimum F_{DR} .

D. Sensitivity

To accurately realize a transfer function using analog integrated circuits, the circuit components of the filter must match closely, and the sensitivity of the transfer function to the values of the filter's components must be low. This sensitivity depends on the filter network, and thus depends on the state-space representation. Hence, the sensitivity of the transfer function of a particular state-space representation is an important criterion for the comparison of different network realizations.

The sensitivity of the transfer function $H(s)$ to the component values of a state-space realization depends on the sensitivity to the entries of the state matrices A , B and C . Here we will consider an absolute sensitivity measurement, which can be used to establish the relationship between the absolute changes Δa_{ij} , Δb_i and Δc_j [43].

The absolute sensitivities of $H(s)$ with respect to variations of the coefficients are given by [39]

$$\begin{aligned}
 S_{a_{ij}}(s) &= \frac{\partial H(s)}{\partial a_{ij}} = c(sI - A)^{-1} e_i e_j (sI - A)^{-1} b \\
 &= S_{b_i}(s) S_{c_j}(s) \\
 S_{b_i}(s) &= \frac{\partial H(s)}{\partial b_i} = c(sI - A)^{-1} e_i \\
 S_{c_i}(s) &= \frac{\partial H(s)}{\partial c_i} = e_i (sI - A)^{-1} b
 \end{aligned} \tag{31}$$

where e_i is the unit vector with i th element unity. In the case of a statistical derivation analysis, the frequency dependent variances of the transfer function are defined as

$$\begin{aligned}
 \Sigma_{\Delta H, A}^2(s) &= E \left\{ \left| \sum_{i=1}^n \sum_{j=1}^n S_{a_{ij}} \Delta a_{ij} \right|^2 \right\} \\
 &= \sum_{i=1}^n \sum_{j=1}^n |S_{a_{ij}}|^2 \\
 \Sigma_{\Delta H, B}^2(s) &= E \left\{ \left| \sum_{i=1}^n S_{b_i} \Delta b_i \right|^2 \right\} \\
 &= \sum_{i=1}^n |S_{b_i}|^2 \\
 \Sigma_{\Delta H, C}^2(s) &= E \left\{ \left| \sum_{j=1}^n S_{c_j} \Delta c_j \right|^2 \right\} \\
 &= \sum_{j=1}^n |S_{c_j}|^2
 \end{aligned} \tag{32}$$

where the variations of the coefficients have been considered statistically independent. Integrating the transfer function variances over the whole frequency range, new sensitivity measures related to the observability and controllability gramians have been defined in [39], and are given by

$$\begin{aligned}
 m_A &\leq \text{Tr}(K) \text{Tr}(W) \\
 m_B &= \text{Tr}(W) \\
 m_C &= \text{Tr}(K)
 \end{aligned} \tag{33}$$

where m_A , m_B , and m_C represent the sensitivity of the matrices A , B , and C to their coefficients, respectively. Finally, the total sensitivity measure, m_T , of a transfer function with respect to the state-space representation matrices can be obtained as

$$m_T = m_A + m_B + m_C \leq \text{Tr}(K) \text{Tr}(W) + \text{Tr}(K) + \text{Tr}(W). \tag{34}$$

The maximum sensitivity measure (worst case sensitivity), considering $m_T = \text{Tr}(K) \text{Tr}(W) + \text{Tr}(K) + \text{Tr}(W)$, can be seen in Fig. 7 as function for the order n of the Morlet wavelet filter, for different state-space representations. One can note that the optimal DR state-space representation will also be optimal with respect to sensitivity. The Schur and Hessenberg forms decomposed from an optimal system also present an optimal




Fig. 7. Sensitivity versus order of the Morlet wavelet filter for different state-space representations.

sensitivity because the orthogonal transformation involved does not affect the sensitivity. The orthonormal ladder structure has a reasonably low sensitivity to coefficient mismatch, close to the optimal case. Both (controllable and observable) canonical forms have the worst sensitivity measures for high order filters, as expected.

E. Dynamic Range-Sparsity-Sensitivity Figure-of-Merit (DRSS)

To conclude the discussion on state-space descriptions we will present a new figure of merit that expresses the correlation of dynamic range, sparsity, and sensitivity parameters. In order to relate these three aspects, we introduce the dynamic range-sparsity-sensitivity figure-of-merit (DRSS), given by

$$\text{DRSS} = \frac{N_c \times F_{\text{DR}} \times m_T}{n}. \quad (35)$$

As one can see from (35), the objective functional, F_{DR} , gives the relative deterioration of the dynamic range, N_c/n is related to the sparsity of the system and m_T defines the sensitivity of matrices A , B , and C with respect to their coefficients. Fig. 8 shows the DRSS figure-of-merit versus the order of the system. Again, the orthonormal ladder structure presents the best performance (minimum DRSS) compared to the other state-space descriptions. Thus, we can state that the orthonormal ladder structure is the optimal state-space description for a system design where the most important requirements are the dynamic range, sparsity, and sensitivity. Although the state-space description obtained from the optimization procedure in Subsection A gives us the optimal representation with respect to dynamic range and sensitivity, its matrices are fully dense, and consequently, its SFOM and DRSS performances are relatively poor.

In this section, we presented the description and a comparison of several state-space representations. The analysis was based on dynamic range, sparsity and sensitivity properties, which are the most relevant aspects for an ultra low-power analog dynamic system. From the two new figures-of-merit described above, *viz.*, the SFOM and DRSS, we concluded that the orthonormal ladder structure [42] is the optimal state-space representation, and, therefore, will be used for the design of continuous-time analog wavelet filters in the next section.




Fig. 8. DRSS figure-of-merit versus order of the Morlet wavelet filter for different state-space representations.

V. ULTRA LOW-POWER CONTINUOUS-TIME ANALOG WAVELET FILTER DESIGN

An n th-order linear differential equation, which describes a filter of the same order, can be implemented by means of n intercoupled integrators. An (on-chip) integrator is an electronic circuit that realizes the transfer function G/sC , where G is implemented by a (trans)conductor (which converts voltage into current), and the integrating component $1/sC$ is usually realized by a (trans)capacitance (integrating current into voltage). Therefore, integrators can be seen as the main building blocks of filter topologies, and, consequently, we will use an ultra-low-power integrator topology to implement the desired wavelet filter.

The trend towards ultra low-power integrated continuous-time filters has increased the interest in new design techniques for analog integrated filters. The current state-of-the-art design approaches for such filters are transconductance amplifier-capacitor (G_m -C) and dynamic translinear (log-domain) methods. In the field of medical electronics, active filters with large time constants are often required to attain low cutoff-frequencies, in the Hz and sub-Hz ranges. In order to avoid large capacitor values on chip, very low transconductances should be implemented to design very low cutoff-frequency filters. Therefore, a G_m -C structure is a natural choice to implement the integrators, as long as very-low transconductance values, typically a few nA/V or less, can be achieved.

The filter design example that follows is based on an orthonormal ladder structure and employs the nA/V transconductor described in [27]. It is based on the use of CMOS transistors operating in the strong-inversion triode region (SI-TR). As shown in [44], transistors kept in SI-TR benefit from a lower G_m/I_D ratio than the ones operating in saturation (active) or weak-inversion regions. This means that, for a particular bias current I_D , the SI-TR transconductor presents the lowest G_m value. In addition, triode-based transconductors have better linearity performance than transconductors with transistors operating in saturation.

The SI-TR transconductor is shown in Fig. 9. It should be stated here that in this circuit not all transistors are in the triode region. Only M_{1A} and M_{1B} are biased in the strong-inversion triode region. M_{2A} and M_{2B} operate in weak inversion triode.

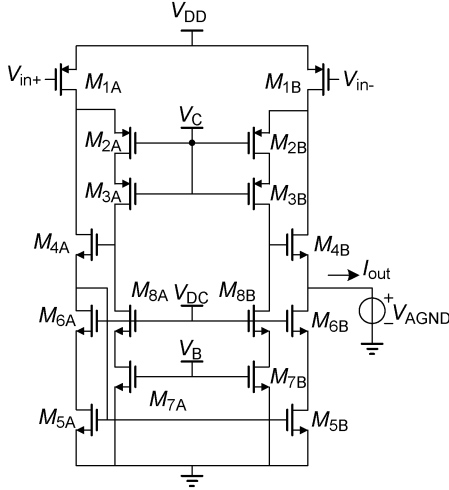


Fig. 9. Schematic of the nA/V CMOS triode-transconductor [27].

The other transistors are operating in the weak inversion saturation region.

Input transistors $M_{1A} - M_{1B}$ have their drain voltages regulated by an auxiliary amplifier that comprises the current conveyor $M_{2A} - M_{2B}$, $M_{3A} - M_{3B}$ and $M_{4A} - M_{4B}$, and bias current sources $M_{7A} - M_{7B}$ and $M_{8A} - M_{8B}$. A low-voltage cascode current mirror comprising $M_{5A} - M_{5B}$ and $M_{6A} - M_{6B}$ provides a single-ended output. Internal voltages V_B , V_C , and V_{DC} are derived from the bias circuit shown in Fig. 10, which is based on the one presented in [44]. The bias generator is structurally alike the transconductor so that the external voltage V_{TUNE} is reflected onto the drain of $M_{1A} - M_{1B}$. Referring V_{TUNE} to V_{DD} , one can define the G_m of the transconductor as

$$G_m = \beta \cdot V_{TUNE} \quad (36)$$

with $\beta = (W/L)\mu C_{ox}$. In order to obtain a G_m -C filter realization using the proposed integrator, we must be able to map the corresponding filter coefficients on the respective G_m values. From the transconductance definition in (36), one can notice that we may vary the value of G_m by changing the drain-source voltage (V_{TUNE}) or, alternatively, β (by the aspect ratio W/L) of transistors M_{1A} and M_{1B} .

In Fig. 11, the Gaussian wavelet (gaus1) has been approximated using the L_2 approach, with the impulse response of the sixth-order transfer function

$$H(s) = \frac{-0.08946s^5 - 0.1683s^4 - 8.326s^3 + 6.642s^2 - 139.6s}{s^6 + 5.927s^5 + 30.52s^4 + 83.11s^3 + 163.6s^2 + 176.6s + 93.29} \quad (37)$$

The time shift involved equals 2.0, with a corresponding L_2 truncation error of 0.0238. The L_2 approximation error of the (time-reversed, shifted) truncated wavelet function by the impulse response is 0.0132, yielding a total L_2 error of 0.0272.

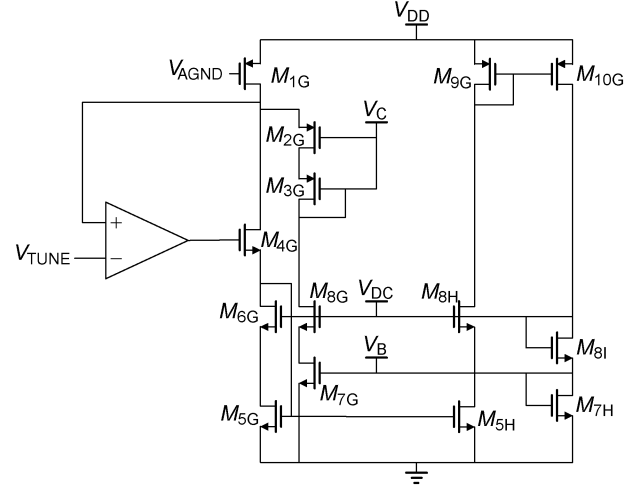



Fig. 10. Circuit diagram of the bias generator.

Fig. 11. Sixth-order L_2 approximation of the gaus1 wavelet with shift $t_0 = 2.0$.

Since the wavelet is normalized to unit energy, this can be regarded as a maximal 2.72% error for the wavelet transform; see also (7). Using the state-space realization method of [42] the following matrices were obtained, as shown in (38) at the bottom of the next page.

To implement this filter, we chose to scale C by a factor of 5 (to bring all nonzero entries into a comparable numerical range) and we set the last two entries of C to zero (to reduce the total number of transistors and thus chip area and power consumption required for the implementation). This increases the total L_2 error to 4.07%.

The block diagram of the wavelet filter is presented in Fig. 12. The values of G_m and the total capacitance required to implement the transfer function in (37) are also shown in Fig. 12. The sensitivity of the total L_2 error of the wavelet filter transfer function on each of the individual values of G_m used in this implementation, is shown in Table III. There, the total L_2 error is shown (as a percentage) which results when different scaling factors are applied to the individual G_m values. We conclude that for small inaccuracies the total L_2 error hardly increases. In practice, the impact of such deviations on the wavelet transform

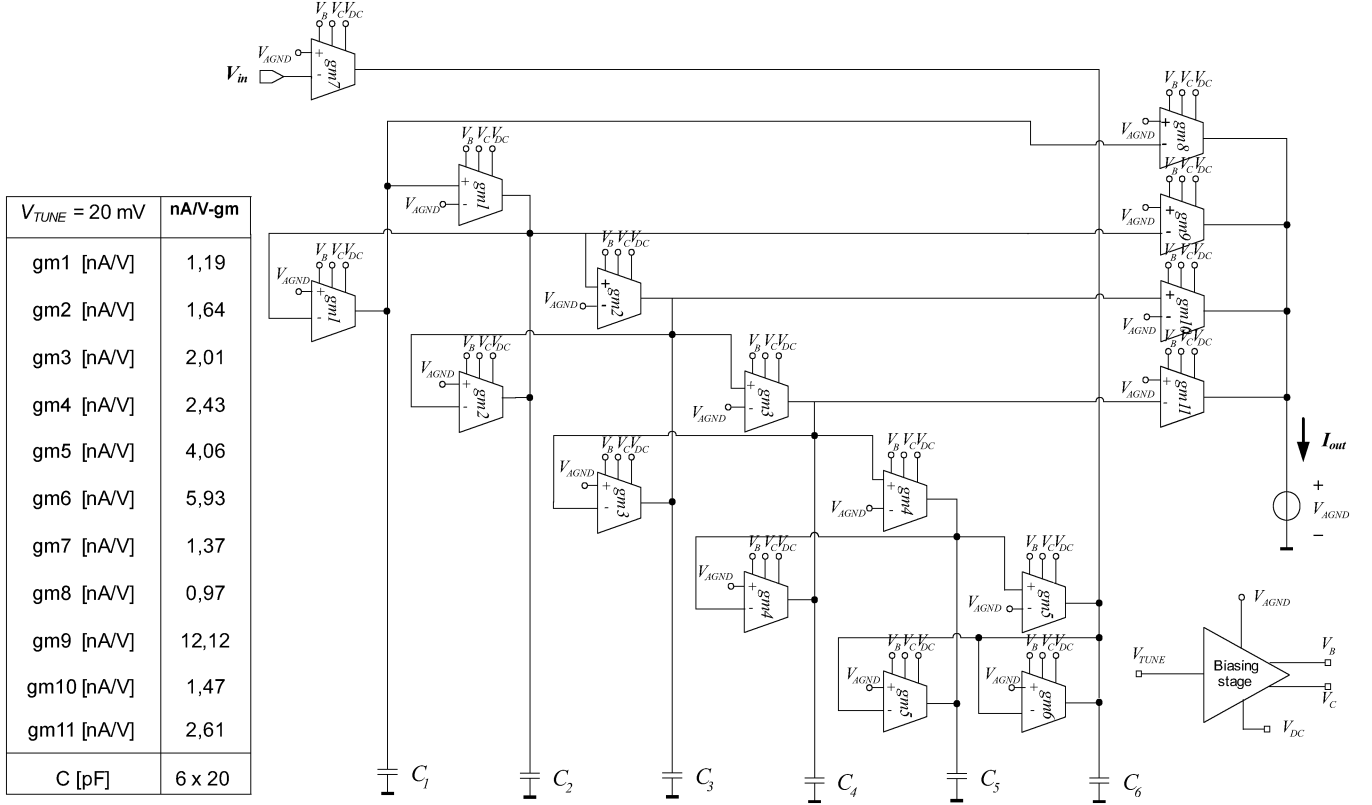


Fig. 12. Block diagram of the wavelet filter.

TABLE III

L_2 ERROR OF THE WAVELET IMPLEMENTATION FOR INDIVIDUAL SCALING OF THE COEFFICIENTS G_m BY FACTORS 0.95, 0.98, 0.99, 1.01, 1.02 AND 1.05. THE L_2 ERROR OF THE REFERENCE IMPLEMENTATION IS 4.07%

	0.95	0.98	0.99	1.01	1.02	1.05
G_{m1}	7.65%	4.47%	4.00%	4.65%	5.60%	9.37%
G_{m2}	9.17%	5.00%	4.20%	4.64%	5.70%	10.0%
G_{m3}	7.57%	4.82%	4.28%	4.22%	4.68%	7.15%
G_{m4}	6.90%	4.50%	4.12%	4.34%	4.88%	7.33%
G_{m5}	6.87%	4.61%	4.21%	4.21%	4.61%	6.68%
G_{m6}	5.50%	4.32%	4.14%	4.13%	4.29%	5.29%
G_{m7}	6.35%	4.48%	4.16%	4.22%	4.59%	6.54%
G_{m8}	4.12%	4.09%	4.08%	4.06%	4.06%	4.05%
G_{m9}	6.11%	4.39%	4.12%	4.25%	4.62%	6.53%
G_{m10}	4.17%	4.10%	4.08%	4.06%	4.05%	4.05%
G_{m11}	4.27%	4.12%	4.09%	4.06%	4.06%	4.13%

will also depend on the signal being processed and is usually smaller.

The wavelet G_m -C filter has been simulated using AMIS 0.35- μm CMOS transistor models. Fig. 13 shows the simulated and L_2 approximated impulse responses of the wavelet filter for $V_{TUNE} = 20 \text{ mV}$. An excellent approximation to the L_2 approximated first Gaussian wavelet (gaus1) confirms the performance of the G_m -C filter.


Table IV presents the simulated filter specifications. The total filter current consumption ranges from 25 nA to 183 nA, depending on the scale, while operating from a 1.8-V supply. The figure-of-merit (FOM) commonly used for comparison of different filters and defined as $\text{FOM} = P/(n \times f_c \times \text{DR})$, in which P is the total power consumption of the wavelet filter, n is the order of the filter and DR is the filter's dynamic range [for a 1% total harmonic distortion (THD)], has been given as well. It must be noted though that, although the filter approximation and filter topology have been optimized for power consumption and dynamic range, due to the inherently larger amount of noise generated in the triode transconductor, the FOM is relatively large.

$$A = \begin{pmatrix} 0 & 1.185 & 0 & 0 & 0 & 0 \\ -1.185 & 0 & 1.637 & 0 & 0 & 0 \\ 0 & -1.637 & 0 & 2.007 & 0 & 0 \\ 0 & 0 & -2.007 & 0 & 2.431 & 0 \\ 0 & 0 & 0 & -2.431 & 0 & 4.062 \\ 0 & 0 & 0 & 0 & -4.062 & -5.927 \end{pmatrix}$$

$$B = (0 \ 0 \ 0 \ 0 \ 0 \ 1.374)^T$$

$$C = (-0.1948 \ -2.424 \ 0.2941 \ -0.5214 \ -0.03016 \ -0.06514)$$

(38)

Fig. 13. Simulated and L_2 -approximated impulse response.TABLE IV
FILTER SPECIFICATIONS

$V_{DD}=1.8V$, V_{TUNE} [mV]	10	20	40	80
Supply current drawn [nA]	25	49.5	96.4	182.6
Power dissipation [nW]	45	89.1	173.5	328.7
Bandwidth [Hz]	6.2	12.5	25.5	49.7
Center frequency [Hz]	5.9	11.8	22.9	45.7
DR (THD=1%) [dB]	16	20	25	28
V_{in} (THD=1%) [mV _{RMS}]	173	179	241	236
FOM	7.7E-11	6.3E-11	5E-11	4.3E-11
Supply voltage range [V]	1.5 - 2.1			


Better DRs and FOMs are possible using transconductors or integrator circuits based on the use of transistors that operate in weak-inversion saturation. This is, however, beyond the scope of this paper.

Finally, in order to implement a wavelet transform, we need to be able to scale and shift in time (and, consequently in frequency) the *gaus1* function. By changing the values of V_{TUNE} accordingly we implement different scales, while preserving the impulse response waveform, as one can see in Fig. 14(a). Fig. 14(b) illustrates the frequency response (magnitude only) for 4 dyadic scales with center frequencies ranging from 6 Hz to 46 Hz for V_{TUNE} varying from 10 mV to 80 mV, respectively.

VI. CONCLUSION

The L_2 approximation approach for fitting the impulse response of a linear system performs markedly better than earlier Padé based methods, in particular at the beginning of the signal. The L_2 criterion offers a conceptual appeal for the problem at hand. With the current methodology a wide range of continuous-time wavelets can be implemented in an uniform manner. For the optimization a local search technique is employed and a procedure has been presented to find a suitable starting point that helps to avoid that the optimization terminates in a poor local optimum. Further use of randomness in the vicinity of this starting point may further decrease the chance of finding a sub-optimal solution.

After a suitable low-order linear system has been designed, state-space optimization is employed to optimize the dynamic

Fig. 14. Wavelet filter scaling by changing V_{TUNE} . (a) Impulse response. (b) Frequency response (magnitude only).

range of the system, without sacrificing too much sparsity and guarding the sensitivity of the system in the process. In order to come to a well balanced tradeoff the dynamic range-sparsity-sensitivity figure-of-merit is introduced. From this we deduced that the orthonormal ladder structure is the best candidate for the design of continuous-time analog wavelet filter topologies.

Finally, to illustrate the design procedure presented here, an L_2 -approximated orthonormal “*gaus1*” wavelet filter using G_m -C integrator building blocks was presented. Simulation results on an actual circuit design confirm that indeed various scales of the wavelet transform can be implemented in an analog fashion consuming little power.

ACKNOWLEDGMENT

The authors would like to thank STW and the users committee for their support.

REFERENCES

- [1] G. Strang and T. Nguyen, *Wavelets and Filter Banks*. Cambridge, U.K.: Wellesley-Cambridge Press, 1996.
- [2] S. Mallat, *A Wavelet Tour of Signal Processing*. New York: Academic, 1999.

- [3] D. A. Sherman, "An introduction to wavelets with electrocardiology applications," *Herzschr Elektrophys.*, vol. 9, pp. 42–52, 1998.
- [4] S. Kadambe, R. Murray, and G. F. Boudreaux-Bartels, "Wavelet transform-based QRS complex detector," *IEEE Trans. Biomed. Eng.*, vol. 46, no. 7, pp. 838–848, Jul. 1999.
- [5] C. Li, C. Zheng, and C. Tai, "Detection of ECG characteristic points using wavelet transforms," *IEEE Trans. Biomed. Eng.*, vol. 42, no. 1, pp. 21–28, Jan. 1995.
- [6] J. Sahambi, S. Tandon, and R. Bhatt, "Using wavelet transforms for ECG characterization," *IEEE Eng. Med. Biol. Mag.*, vol. 16, no. 1, pp. 77–83, 1997.
- [7] P. E. Tikkanen, "Nonlinear wavelet and wavelet packet denoising of electrocardiogram signal," *Biol. Cybern.*, vol. 80, no. 4, pp. 259–267, Apr. 1999.
- [8] F. Mochimaru, Y. Fujimoto, and Y. Ishikawa, "Detecting the fetal electrocardiogram by wavelet theory-based methods," *Progress in Biomed. Res.*, vol. 7, no. 3, pp. 185–193, Sep. 2002.
- [9] K. G. Oweiss, A. Mason, Y. Suhail, A. M. Kamboh, and K. E. Thomson, "A scalable wavelet transform VLSI architecture for real-time signal processing in high-density intra-cortical implants," *IEEE Trans. Circuits Syst. I*, vol. 54, no. 6, pp. 1266–1278, Jun. 2007.
- [10] S. A. P. Haddad, R. Houben, and W. A. Serdijn, "Analog wavelet transform employing dynamic translinear circuits for cardiac signal characterization," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, May 2003, vol. 1, pp. 121–124.
- [11] J. M. H. Karel, R. L. M. Peeters, R. L. Westra, S. A. P. Haddad, and W. A. Serdijn, "Wavelet approximation for implementation in dynamic translinear circuits," in *Proc. 16th IFAC World Congr.*, 2005.
- [12] M. Liu, Q. Hu, and Y. He, "A novel analog VLSI implementation of wavelet transform based on SI circuits," in *Proc. Congr. Image Signal Process. (CISP)*, May 2008, vol. 1, pp. 317–323.
- [13] W. Zhao, Y.-G. He, J. Huang, Y. Y. Xie, and Y. Zhang, "Analogue VLSI implementations of wavelet transform based on switched-current technology," in *Proc. Int. Conf. Wavelet Anal. Pattern Recognit. (ICWAPR)*, Nov. 2007, vol. 4, pp. 1787–1792.
- [14] A. J. Casson, D. C. Yates, S. Patel, and E. Rodriguez-Villegas, "An analogue bandpass filter realisation of the continuous wavelet transform," in *Proc. IEEE Int. Conf. Eng. Med. Biol. Soc. (EMBS)*, Aug. 2007, pp. 1850–1854.
- [15] A. J. Casson and E. Rodriguez-Villegas, "An inverse filter realisation of a single scale inverse continuous wavelet transform," in *Proc. IEEE Int. Symp. Circuit Syst. (ISCAS)*, May 2008, pp. 904–907.
- [16] M. Tuckwell and C. Papavassiliou, "An analog Gabor transform using sub-threshold 180-nm CMOS devices," *IEEE Trans. Circuits Syst. I*, vol. 56, no. 12, pp. 2597–2608, Dec. 2009.
- [17] W. Zhao, Y. Sun, and Y. He, "Minimum component high frequency Gm-C wavelet filters based on Maclaurin series and multiple loop feedback," *Electron. Lett.*, vol. 46, no. 1, pp. 34–36, Jan. 2010.
- [18] S. A. P. Haddad and W. A. Serdijn, "Mapping the wavelet transform onto silicon: The dynamic translinear approach," in *Proc. IEEE Int. Symp. Circuit Syst. (ISCAS)*, May 2002, pp. 621–624.
- [19] J. M. H. Karel, R. L. M. Peeters, R. L. Westra, S. A. P. Haddad, and W. A. Serdijn, "An L_2 -based approach for wavelet approximation," in *Proc. CDC-ECC*, 2005.
- [20] J. M. H. Karel, "A wavelet approach to cardiac signal processing for lowpower hardware applications," Ph.D. dissertation, Maastricht Univ., Maastricht, The Netherlands, Dec. 2009.
- [21] S. A. P. Haddad and W. A. Serdijn, *Ultra Low-Power Biomedical Signal Processing: An Analog Wavelet Filter Approach for Pacemakers*. New York: Springer, Dec. 2006.
- [22] S. A. P. Haddad, N. Verwaal, R. Houben, and W. A. Serdijn, "Optimized dynamic translinear implementation of the gaussian wavelet transform," in *Proc. IEEE Int. Symp. Circuit Syst. (ISCAS)*, May 2004, vol. 1, pp. 145–148.
- [23] D. P. W. M. Rocha, "Optimal design of analogue low-power systems, a strongly directional hearing-aid adapter," Ph.D. dissertation, Delft Univ. of Technol., Delft, The Netherlands, Apr. 2003.
- [24] S. A. P. Haddad, J. M. H. Karel, R. L. M. Peeters, R. L. Westra, and W. Serdijn, "Analog complex wavelet filters," in *Proc. IEEE Int. Symp. Circuit Syst. (ISCAS)*, May 2005, pp. 3287–3290.
- [25] S. A. P. Haddad, S. Bagga, and W. A. Serdijn, "Log-domain wavelet bases," *IEEE Trans. Circuits Syst. I*, vol. 52, no. 10, pp. 2023–2032, Oct. 2005.
- [26] S. A. P. Haddad, J. M. H. Karel, R. L. M. Peeters, R. L. Westra, and W. A. Serdijn, "Ultra low-power analog Morlet wavelet filter in 0.18 μm BiCMOS technology," in *Proc. ESSCIRC*, Grenoble, France, Sep. 2005.
- [27] P. R. Agostinho, S. A. P. Haddad, J. A. D. Lima, and W. A. Serdijn, "An ultra low power CMOS pA/V transconductor and its application to wavelet filters," *Analog Integr. Circuits Signal Process.*, vol. 57, no. 1–2, pp. 19–27, Nov. 2008.
- [28] G. Baker, Jr., *Essentials of Padé Approximants*. New York: Academic, 1975.
- [29] A. Bultheel and M. van Barel, "Padé techniques for model reduction in linear system theory: A survey," *J. Comput. Appl. Math.*, vol. 14, pp. 401–438, 1986.
- [30] H. Kamada and N. Aoshima, "Analog Gabor transform filter with complex first order system," in *Proc. SICE*, 1997, pp. 925–930.
- [31] N. Aoshima, "Analog realization of complex first order system and its application to vowel recognition," in *Proc. SICE*, 1995, pp. 1239–1244.
- [32] T. Kailath, *Linear Systems*. Englewood Cliffs: Prentice Hall, 1980.
- [33] B. C. Moore, "Principal component analysis in linear systems: Controllability, observability, and model reduction," *IEEE Trans. Automat. Control*, vol. 26, no. 1, pp. 17–32, Feb. 1981.
- [34] L. Pernebo and L. M. Silverman, "Model reduction via balanced state space representations," *IEEE Trans. Autom. Control*, vol. 27, no. 2, pp. 382–387, Apr. 1982.
- [35] A. Kapela, R. D. Berger, A. Achim, and A. Bezerianos, "Wavelet variance analysis of high-resolution ECG in patients prone to VT/VF during cardiac electrophysiology studies," in *Proc. 14th Int. Conf. Digital Signal Process. (DSP 2002)*, Jul. 2002, vol. 2, pp. 1133–1136.
- [36] P. S. Addison, J. N. Watson, G. R. Clegg, M. Holzer, and C. E. R. F. Sterz, "Evaluating arrhythmias in ECG signals using wavelet transforms," *IEEE Eng. Med. Biol. Mag.*, vol. 19, no. 5, pp. 104–109, Sep.–Oct. 2000.
- [37] P. S. Addison, "Wavelet transforms and the ECG: A review," *Physiol. Meas.*, vol. 26, pp. R155–R199, 2005.
- [38] M. J. Burke and M. Nasor, "Wavelet based analysis and characterization of the ECG signal," *J. Med. Eng. Technol.*, vol. 28, no. 2, pp. 47–55, Mar.–Apr. 2004.
- [39] L. Thiele, "On the sensitivity of linear state-space systems," *IEEE Trans. Circuits Syst.*, vol. 33, no. 5, pp. 502–510, May 1986.
- [40] W. M. Snelgrove and A. S. Sedra, "Synthesis and analysis of state-space active filters using intermediate transfer function," *IEEE Trans. Circuits Syst.*, vol. 33, no. 3, pp. 287–301, Mar. 1986.
- [41] G. Groenewold, "Optimal dynamic range integrators," *IEEE Trans. Circuits Syst. I*, vol. 39, no. 8, pp. 614–627, Aug. 1992.
- [42] W. Snelgrove, D. A. Johns, and A. Sedra, "Orthonormal ladder filters," *IEEE Trans. Circuits Syst.*, vol. 36, no. 3, pp. 337–343, Mar. 1989.
- [43] W. K. Chen, *The Circuits and Filters Handbook*. Boca Raton, FL: CRC, IEEE Press, 1995.
- [44] J. A. de Lima and W. A. Serdijn, "A compact nA/V triode-MOSFET transconductor," *Electron. Lett.*, vol. 41, no. 20, pp. 1113–1114, Sep. 2005.



Joël M. H. Karel (S'04–M'08) was born in 1978. He received the "Kandidaat Informatica" degree in computer science from Hasselt University, Hasselt, Belgium, in 1999. In the same year he attended a special four week track on knowledge engineering and computer science at Baylor University, Waco, TX. He received the Doctoraal (M.Sc. degree) in knowledge engineering from the Maastricht University, Maastricht, The Netherlands, in 2001, after completing his thesis on modeling atrial fibrillation at the Medtronic Bakken Research Center. He then obtained a joint position at the Maastricht University and MaTeUM b.v. on simulation. He started the Ph.D. degree in the field of biomedical signal processing and systems theory with application for cardiac pacemakers in 2004 at the Department of Mathematics on the BioSens project, Maastricht University. He received the Ph.D. degree after defending his thesis titled "A wavelet approach for cardiac signal processing for low-power hardware applications" from Maastricht University in 2009.

Since 2007, he has combined his Ph.D. research with a part-time position as an Assistant Professor and since 2008 he has been a full-time Assistant Professor in the Department of Knowledge Engineering, Maastricht University. His research interests are in the field of biomedical engineering, signal processing, wavelet theory and linear systems theory.



Sandro A. P. Haddad (M'11) was born in Anápolis, Brazil, on February 8, 1977. He received the B.S. degree (with honors) in electrical engineering from the University of Brasilia (UnB), Brasilia, Brazil, in 2000 and the Ph.D. degree from the Electronics Research Laboratory, Delft University of Technology (TUDelft), Delft, The Netherlands, in 2006 with the thesis entitled "Ultra Low-Power Biomedical Signal Processing: An Analog Wavelet Filter Approach for Pacemakers." His project was part of BioSens (Biomedical Signal Processing Platform for

Low-Power Real-Time Sensing of Cardiac Signals).

From 2007 to 2010, he worked at Freescale Semiconductor as an Analog IC Designer. In June 2010, he joined the University of Brasilia (UnB) as an Associate Professor. His research interests include low-voltage, ultra-low-power analog electronics and biomedical systems, and high-frequency analog integrated circuits for UWB communications. He authored and coauthored more than 40 publications.

Dr. Haddad was awarded the best graduated student of the year, from the Faculty of Electrical Engineering, UnB.



Senad Hiseni (S'10) was born in Bosnia and Herzegovina, in 1978. He received the Ing. degree in electrical engineering from the Haagse Hogeschool, The Hague, The Netherlands, and the M.Sc. degree from the Delft University of Technology, Delft, The Netherlands. He is currently working towards the Ph.D. degree at the Department of Microelectronics at the Faculty of Electrical Engineering, Mathematics, and Computer Science, Delft University of Technology.

His research interests include low-voltage and ultra-low-power analog integrated circuits for biomedical applications.



Ronald L. Westra was born in Amsterdam, The Netherlands, in 1956. He received the B.Sc. degrees in theoretical physics and astronomy and M.Sc. degrees in theoretical physics and mathematics from Utrecht University, Utrecht, The Netherlands, and the Ph.D. degree in computer vision and image analysis from the Maastricht University, Maastricht, The Netherlands.

He held various positions as a Researcher in theoretical physics and mathematics in science and industry, including the University of Amsterdam (UvA) and the Free University of Brussels (ULB). He currently works at the Department of Knowledge Engineering, Maastricht University, as an Associate Professor and Head of the Biomathematics and Bioinformatics Research Group. He is involved in research in biomedical signal analysis and mathematical systems biology.



Wouter A. Serdijn (M'98–SM'08–F'11) was born in Zoetermeer ("Sweet Lake City"), The Netherlands, in 1966. He received the M.Sc. (*cum laude*) and Ph.D. degrees from Delft University of Technology, Delft, The Netherlands, in 1989 and 1994, respectively.

His research interests include low-voltage, ultra-low-power and ultra wideband analog integrated circuits for wireless communications, pacemakers, cochlear implants, portable, wearable, implantable, and injectable ExG recorders, and neurostimulators. He is coeditor and coauthor of the

books *Ultra Low-Power Biomedical Signal Processing: An Analog Wavelet Filter Approach for Pacemakers* (Springer, 2009), *Circuits and Systems for Future Generations of Wireless Communications* (Springer, 2009), *Power Aware Architecting for Data Dominated Applications* (Springer, 2007), *Adaptive Low-Power Circuits for Wireless Communications* (Springer, 2006), *Research Perspectives on Dynamic Translinear and Log-Domain Circuits* (Kluwer, 2000), *Dynamic Translinear and Log-Domain Circuits* (Kluwer, 1998), and *Low-Voltage Low-Power Analog Integrated Circuits* (Kluwer, 1995). He authored and coauthored six book chapters and more than 200 publications and presentations. He teaches analog electronics, analog signal processing, micropower analog IC design and electronic design techniques.

Dr. Serdijn received the Electrical Engineering Best Teacher Award in 2001 and 2004. He has served as an Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS-I: REGULAR PAPERS (2004–2005) and the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS-II: EXPRESS BRIEFS (2002–2003 and 2006–2007), as Deputy Editor-in-Chief for IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS-I: REGULAR PAPERS (2008–2009), as member of the Editorial Board of *Analog Integrated Circuits and Signal Processing* (Springer), as member of the Editorial Board of the *Journal on Low Power Electronics*, as Tutorial Session co-chair for ISCAS 2003, as Analog Signal Processing Track Co-Chair for ISCAS 2004 and ISCAS 2005, as Analog Signal Processing Track Chair for ICECS 2004, as Technical Program Committee member for the 2004 International Workshop on Biomedical Circuits and Systems, as International Program Committee member for IASTED CSS 2005 and CSS 2006, as Technical Program Committee member for APCCAS 2006, as Technical Program Committee member for the IEEE Biomedical Circuits and Systems Conference (BioCAS 2006, BioCAS 2007 and BioCAS 2008), as Special-Session Chair for ISCAS 2007, as International Program Committee member of the 2009 International Conference on Biomedical Electronics and Devices, as Special Session Chair for ISCAS 2009, as Special Sessions Chair for ICECS 2009, as Technical Program Committee member for ICUWB 2009, as Technical Program Chair for ISCAS 2010, as Technical Program Chair for BioCAS 2010, as chair of the Analog Signal Processing Technical Committee of the IEEE Circuits and Systems society, as a member of the CAS-S Long Term Strategy Committee and as a member of the CAS-S Board of Governors Nominations Committee. He currently serves as a member of the Board of Governors (BoG) of the Circuits and Systems Society (second term), a member of the Conference Division of the CAS-S BoG, as Technical Program Co-Chair for ISCAS 2012, as Editor-in-Chief for IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS-I: REGULAR PAPERS (2010–2011) and as a member of the Steering Committee of the IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS (T-BIOCAS). He is a fellow and a mentor of the IEEE.



Ralf L. M. Peeters was born in 1964. He received the M.Sc. degree in applied mathematics from the Delft University of Technology, Delft, The Netherlands, in 1988 and the Ph.D. degree from the Economics Faculty of the Free University, Amsterdam, The Netherlands, in 1994 with a thesis in the area of system identification.

Since 1994, he has been with Maastricht University, Maastricht, The Netherlands, in what is now called the Department of Knowledge Engineering. In 2007, he became Head of the Department and he was appointed Full Professor in 2008. His research interests include: system identification (parameterization of systems, identifiability issues, model order reduction and model approximation); signal processing (wavelet analysis and orthogonal wavelet design, also for 2-D image processing); numerical mathematics (optimization, Lyapunov equations, large eigenvalue problems); computer algebra applications to systems theory (identifiability, global optimization, model approximation); applications of knowledge engineering to medicine and the life sciences (ECG and EEG analysis, medical image analysis, gene regulatory networks, systems biology).