Proceedings of the
44th IEEE Conference on Decision and Control, and
the European Control Conference 2005
Seville, Spain, December 12-15, 2005

An L_2-based approach for wavelet approximation

J.M.H. Karel*, R.L.M. Peeters*, R.L. Westra*, S.A.P. Haddad† and W.A. Serdijn†
*Department of Mathematics, Maastricht University
P.O. Box 616, 6200 MD Maastricht, The Netherlands
Email: {joel.karel, ralf.peeters, westra}@math.unimaas.nl
†Electronics Research Laboratory, Delft University of Technology
Mekelweg 4, 2628 CD Delft, The Netherlands
Email: {w.a.serdijn, s.haddad}@its.tudelft.nl

Abstract— **Computation in the analog domain is very appealing from a power-consumption perspective. To implement a wavelet transform in an analog circuit, the wavelet function can be approximated by a linear system. An approach based on L_2-approximation is presented, that enables largely automated approximation of wavelet functions by impulse responses of linear systems. Various continuous wavelet functions, such as the Gaussian wavelet and Daubechies wavelets of several orders, have been successfully approximated with this approach.**

I. INTRODUCTION

Analog circuits have an advantage over digital circuits due to power-consumption considerations. Particularly A/D (analog to digital) converters, necessary to transform analog sensor information to the digital domain, can be highly power-demanding. Therefore analog circuits are still widely used in applications for which power consumption is a critical issue such as in implantable and portable devices, for instance pacemakers.

In many signal processing applications the wavelet transformation (WT) technique has shown to be extremely useful. The wavelet transform of a signal $f(t)$ over scales σ with a real wavelet $\psi(t)$ is defined as:

$$W_\psi(t,\sigma) = \int_{-\infty}^{\infty} f(\tau)\frac{1}{\sqrt{\sigma}}\psi(\frac{\tau-t}{\sigma})d\tau \qquad (1)$$

The scales σ allow to "zoom in" on interesting parts of the signal. This transformation is usually computed digitally, yielding a high power consumption and poor parallel computational performance for different scales. In [1] however, a method to perform WTs in the analog domain with the technique of Dynamic Translinear (DTL) circuits [2], [3] was introduced. In [1] the Laplace transforms of wavelet functions were approximated by rational functions in the Laplace domain with Padé-approximation. This analog implementation of wavelets yields large power savings over a digital implementation. In [4] it was argued that the Padé-approximation of wavelet functions is not the most suited approach and an alternative approach, based on L_2-approximation that works directly in the time domain, was introduced. A drawback of this approach is that a starting point is required for the application of an iterative local search algorithm to find an optimal approximation. The results have been observed to depend markedly on the choice of starting point due to the existence of local optima. In the current paper the L_2-approximation based approach is discussed and a solution to the problem of finding a good starting point is presented, involving high-order FIR approximation and balance-and-truncate model reduction; see also [5].

The hardware implementation of the wavelets will limit the maximum obtainable accuracy of the wavelet approximation, depending on issues such as the quality of the hardware components and the accuracy of current sources. This consideration will give an upper bound on the required accuracy of the approximation.

II. IMPLEMENTING WAVELET TRANSFORMS IN LINEAR SYSTEMS

Wavelet transforms can be implemented in analog circuits by approximating the *time reversed* wavelet function $\psi(-t)$ with the impulse response $h(t)$ of a linear system $H(s)$, since if a signal $f(t)$ is passed through a linear system, it is convoluted with the impulse response of the system:

$$f(t)*h(t) = \int_{-\infty}^{\infty} f(\tau)h(t-\tau)d\tau \approx \int_{-\infty}^{\infty} f(\tau)\psi(\tau-t)d\tau \quad (2)$$

The wavelet transform $W_\psi(t)$ at scale $\sigma = 1$ on the other hand is defined as the L_2-inner product of the input signal $f(\tau)$ and the shifted wavelet $\psi(\tau-t)$:

$$W_\psi(t) = \langle f(\tau), \psi(\tau-t)\rangle = \int_{-\infty}^{\infty} f(\tau)\psi(\tau-t)d\tau \qquad (3)$$

This shows that if the impulse response $h(t)$ of the linear system approximates the time-reversed wavelet function then the wavelet transformation is approximated by the output of that linear system.

For obvious physical reasons only the hardware implementation of (strictly) causal stable filters is feasible. The effective support of wavelets however is generally located around the origin. As a result the wavelet will have to be *time-shifted* to avoid truncation of a large part of its energy as illustrated in Fig. 1. The selection of the time-shift t_0 can be a delicate process since shifting too much may make

This paper is part of the BioSens project that is being funded by STW (Project no. DTC 6418). web: http://www.math.unimaas.nl/biosens

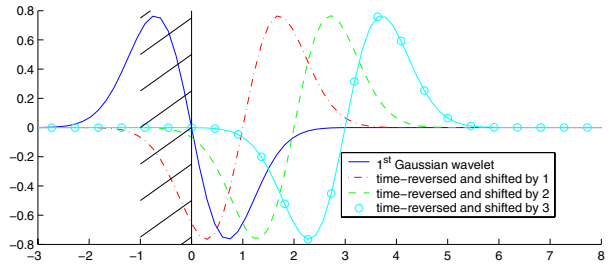

Fig. 1. The importance of time-shifting the Gaussian wavelet

TABLE I

EFFECT OF TIME-SHIFT ON REQUIRED ORDER, NORM OF MISFIT AND ENERGY LOSS

order	norm	shift 2.0	shift 2.5	shift 3.0	shift 3.5
3	L_1-norm	0.997	1.466	1.871	1.980
5	L_1-norm	0.117	0.265	0.496	0.789
7	L_1-norm	0.017	0.027	0.071	0.153
9	L_1-norm	0.016	0.002	0.007	0.021
3	L_2-norm	0.397	0.551	0.678	0.696
5	L_2-norm	0.048	0.101	0.178	0.269
7	L_2-norm	0.007	0.010	0.025	0.053
9	L_2-norm	0.010	0.001	0.003	0.007
3	L_∞-norm	0.490	0.539	0.485	0.430
5	L_∞-norm	0.077	0.151	0.232	0.303
7	L_∞-norm	0.005	0.017	0.041	0.076
9	L_∞-norm	0.002	0.001	0.004	0.011
3	CPU time	0.630	0.820	1.380	1.700
5	CPU time	0.860	0.990	1.240	1.430
7	CPU time	17.010	2.340	1.980	1.820
9	CPU time	933.170	62.390	10.390	5.520
	energy loss	5.5E^{-4}	7.4E^{-6}	3.5E^{-8}	6.1E^{-11}

the approximation of the wavelet hard, due to the existence of a relatively flat slope in the beginning of the support of the wavelet, and induce a large time-delay. Shifting too little yields truncation of energy and may result in an integral that is unequal to zero, thus not stopping the zero frequency. This trade-off between truncation of energy and approximation complexity is illustrated in Table I. In this table various norms and the computation time in seconds from a starting point to an approximation of a Gaussian wavelet for various orders of systems and various time-shifts are listed. Not only the L_2-norm that has been used in the approximation, but also the commonly used L_1-norm and L_∞-norm are shown as a reference. The time-reversed and time-shifted wavelet $\widetilde{\psi}(t)$ now becomes:

$$\widetilde{\psi}(t) = \psi(t_0 - t), \tag{4}$$

and the impulse response $h(t)$ will have to approximate $\widetilde{\psi}(t)$. Note that Table I shows that for the time-shift $t_0 = 2.0$ the approximation with a dynamical system of order 9 involves an L_2-approximation error that is larger than for the approximation of order 7. This is caused by (slower) convergence to a different local optimum. Note also that the CPU times for approximations of order 9 are substantially larger than for lower order approximations. If the time-shift t_0 increases, higher order approximations are required to achieve the same approximation accuracy.

III. L_2-APPROXIMATION OF WAVELET FUNCTIONS

In this section the choice for L_2-approximation, the model and wavelet related issues will be discussed.

A. Motivation

From a conceptual point of view, the L_2-approximation approach is very appropriate for the problem at hand:

- It is appropriate to use the L_2-norm to measure the quality of an approximation $h(t)$ of the wavelet $\widetilde{\psi}(t)$ since $W_\psi(t)$ involves an L_2-inner product.
- It is desirable that the approximation $h(t)$ of $\widetilde{\psi}(t)$ behaves equally well for all time instances t since it will be used as a convolution kernel.
- L_2-approximation allows for a description in the time domain as well as in the Laplace domain due to Parseval's equality.

The first point emphasizes that the L_2-norm is appropriate since it encompasses the very definition of the wavelet transform which is an L_2-inner product. The second point is based on the fact that a continuous wavelet transform will be performed. The convolution kernel will thus be used with an arbitrary shift which makes that the approximation accuracy at each time instance is equally important. The third point is that minimization of $\|\widetilde{\psi} - h\|$ is, due to Parseval's equality, equivalent to the minimization of the L_2-norm of the difference between the Laplace transforms $\widetilde{\Psi}(s)$ and $H(s)$ of $\widetilde{\psi}(t)$ and $h(t)$ respectively, over the imaginary axis $s = i\omega$. Since wavelets are intended to gather information about the signal $f(t)$ both in the time domain and in the frequency domain, it is important to have approximation quality specifications in both domains.

There is one obvious disadvantage to the L_2-approximation approach that we already have come across at the end of the previous section: the existence of *local optima*. There generally cannot be given hard guarantees that an optimum obtained by an iterative local search optimization technique is a global optimum. The local optimum at which the optimization terminates depends on the choice of an initial starting point provided by the user. A good initial point may help to avoid local optima. An automated procedure that selects a starting point that performs well, according to simulation, is described in section IV.

B. Model

From the theory of linear systems (see, e.g. [6]) it is well known that any strictly causal linear filter of finite order n can be represented in the time domain as a state-space system (A,B,C) described by a set of associated differential equations of the form:

$$\dot{x}(t) = Ax(t) + Bu(t), \tag{5}$$
$$y(t) = Cx(t), \tag{6}$$

where $u(t)$ is a scalar input[1] and $y(t)$ a scalar output. The state vector at time t is denoted by $x(t) \in \mathbb{R}^n$. Note that a "direct feed-through" or *dc*-term is not required since strict causality is assumed. The impulse response function $h(t)$ and its Laplace transform $H(s)$, better known as the *transfer function* of the system, are given by:

$$h(t) = Ce^{At}B, \tag{7}$$
$$H(s) = C(sI_n - A)^{-1}B. \tag{8}$$

For the generic situation of stable systems with distinct poles, the impulse response function $h(t)$ is a linear combination of damped exponentials and exponentially damped harmonics. From this it is possible to propose an explicitly parameterized class of impulse response functions that is suitable for approximating a large number of continuous wavelets such as the Gaussian wavelets as described in [4]. For instance, if a 5th order approximation is attempted, this parameterized class of functions $h_p(t)$, involving a real parameter vector $p = \{p_1, \ldots, p_{10}\}$, will typically have the following form:

$$\begin{aligned} h_p(t) = {} & p_1 e^{p_6 t} + p_2 e^{p_7 t}\sin(p_8 t) + p_3 e^{p_7 t}\cos(p_8 t) \\ & + p_4 e^{p_9 t}\sin(p_{10} t) + p_5 e^{p_9 t}\cos(p_{10} t) \end{aligned} \tag{9}$$

The parameters p_6, p_7 and p_9 must be strictly negative for reasons of stability. A corresponding state-space representation is given by:

$$A = \begin{pmatrix} p_6 & 0 & 0 & 0 & 0 \\ 0 & p_7 & p_8 & 0 & 0 \\ 0 & -p_8 & p_7 & 0 & 0 \\ 0 & 0 & 0 & p_9 & p_{10} \\ 0 & 0 & 0 & -p_{10} & p_9 \end{pmatrix}, \quad B = \begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \\ 1 \end{pmatrix},$$
$$C = \begin{pmatrix} p_1 & p_2 & p_3 & p_4 & p_5 \end{pmatrix}. \tag{10}$$

To approximate a wavelet function $\widetilde{\psi}(t)$ the following minimization problem has to be solved:

$$\min_p \left(\sqrt{\int_0^\infty \left(h_p(t) - \widetilde{\psi}(t)\right)^2 dt} \right) \tag{11}$$

for which standard optimization techniques are widely available in the literature and in software packages such as Matlab, Maple, Mathematica, and so on. In the current research the tail of (11) was truncated due to the fast decay, but it was left long enough to ensure that the oscillations damped out. Then the integral was approximated with a Riemann sum, which will give a reasonable good approximation since the integral of the wavelet equals zero (as will be discussed in the next paragraph), and finally the least-squares problem was minimized. However if a higher accuracy is required, another integration scheme can be employed.

C. Ensuring a zero wavelet integral

One common property of a wavelet function $\psi(t)$ that was not discussed so far is that its integral is equal to zero:

$$\int_{-\infty}^{\infty} \psi(t)dt = 0. \tag{12}$$

[1]The input was previously denoted by $f(t)$, but for state-space systems the usual convention $u(t)$ will be used.

If this property is not shared by the approximation $h_p(t)$, this will cause an unwanted bias in the approximation of the wavelet transform. Due to truncation, the time-reversed and time-shifted wavelet $\widetilde{\psi}(t)$ is likely not to have a zero integral. Therefore, when ensuring that the approximation $h_p(t)$ satisfies $\int_0^\infty h_p(t)dt = 0$, a certain approximation error becomes unavoidable. Experiments have shown however that the wavelet transform $W_h(t,\sigma)$ obtained with an approximated wavelet $h_p(t)$ having a zero integral is usually closer to the true wavelet transform $W_\psi(t,\sigma)$ than when one uses the truncated wavelet $\widetilde{\psi}$. The approximation $h_p(t)$ does not carry the bias that is introduced by an integral that is unequal to zero, and it obeys the admissibility conditions [7] that define the wavelet framework and enforce desirable properties such as completeness and energy preservation.

For linear filters, the property that the integral of the impulse response function $h(t)$ is zero is equivalent to the property that the step response of the filter tends to zero for large t. Keeping in mind that the impulse response is the derivative of the step response, from the properties of the Laplace transform it holds that

$$\lim_{t\to\infty} \int_0^t h(\tau)d\tau = \lim_{s\downarrow 0} H(s) \tag{13}$$

Therefore, the desired property comes down to $H(0) = 0$. In terms of a state-space representation (A,B,C) we have that

$$H(0) = -CA^{-1}B. \tag{14}$$

For the parameterized class of 5th order approximations (9) this yields the explicit condition:

$$\frac{p_1}{p_6} + \frac{-p_2 p_8 + p_3 p_7}{p_7^2 + p_8^2} + \frac{-p_4 p_{10} + p_5 p_9}{p_9^2 + p_{10}^2} = 0 \tag{15}$$

If preferable this equation may be used to eliminate one of the variables from the optimization problem, such that the constraints are enforced.

The zero integral $H(0) = 0$ effectively introduces one vanishing moment. Additional vanishing moments can be enforced by requiring that additional derivatives at zero equal zero. In the case of k vanishing moments the transfer function will have the following form:

$$H(s) = s^k \frac{P(s)}{Q(s)}, \tag{16}$$

which can be enforced with constrained L_2-approximation. In general the vanishing moments will be approximated to some extend since the L_2-norm is defined both in the time domain and in the frequency domain.

IV. FINDING A GOOD STARTING POINT

It was mentioned earlier that the choice of the starting point can have a considerable impact on the solution that is found by the optimization routine.

To obtain a good starting point for the L_2-approximation approach, one can initially construct a high-order model directly from the sampled wavelet function and reduce this model to the required order. A number of intermediate steps

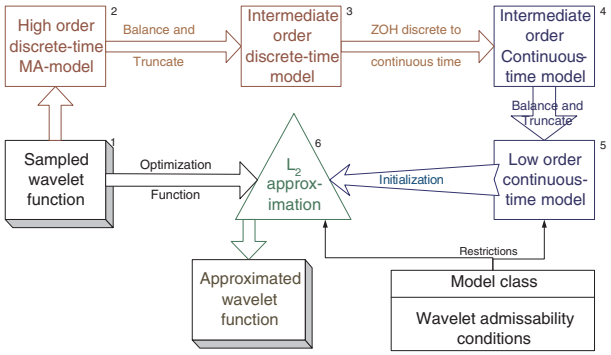

Fig. 2. Obtaining a starting point

are required however, as illustrated in Fig. 2. The approach consists of the following steps that will be further discussed in the remainder of this section:

1) **Sampling the wavelet $\widetilde{\psi}$**
   The time-shifted and reversed wavelet function $\widetilde{\psi}$ is sampled with zero-order-hold, producing a discrete signal g.
2) **Construction of a high order discrete-time FIR-model**
   Signal g is used to construct a high order discrete-time FIR (or moving-average) model, which is represented in state-space form. (typical order 500-2000)
3) **Conversion to an intermediate order discrete-time IIR-model**
   The state-space model is balanced and truncated to yield an accurate reduced order discrete-time model, referred to as an intermediate order model since yet another model reduction step will later take place. (typical order 20-50)
4) **Conversion of the discrete-time IIR-model to a continuous-time IIR-model**
   The discrete-time model is then converted back to continuous-time. Until here, all steps in the procedure have to be performed just once.
5) **Reduction of the continuous-time IIR-model to the desired lower order**
   The intermediate order continuous-time model is reduced to a specified lower order, to be used as a starting point for the optimization technique in the next step. Various reduced orders can be attempted until a satisfactory result is obtained. (typical order 5-12)
6) **L_2-approximation of the wavelet function**
   The low order model obtained in the previous step is used as a starting point for solving the minimization problem (11) under the constraint (14) using an iterative local search optimization technique.

The approach starts with a regularly sampled version of the time-reversed time-shifted wavelet:

$$g_k = \widetilde{\psi}(k \cdot \Delta t), \quad k = 0, \ldots, n-1 \tag{17}$$

where the sampling time Δt and the number of samples n are chosen in accordance with the effective support of the wavelet, keeping in mind that the horizon should be long enough to ensure stability. The sequence $g = \{g_k\}$ is interpreted as a ZOH sampled version of the impulse response of a *continuous-time* system. A corresponding *discrete-time* moving-average system (or finite impulse response (FIR) filter) is constructed from it. It is designed to have the impulse response: $h[0] = 0, \quad h[1] = g_0, \quad h[2] = g_1, \quad \ldots, \quad h[n] = g_{n-1}$. The impulse response of a discrete-time state-space system (A,B,C,D) is given by:

$$h[k] = \begin{cases} D & k = 0 \\ CA^{k-1}B & k > 0 \end{cases} \tag{18}$$

A discrete-time state-space realization $(A_{\mathrm{dh}}, B_{\mathrm{dh}}, C_{\mathrm{dh}}, D_{\mathrm{dh}})$ in controllable companion form, which achieves the finite impulse response h, attains the following form:

$$A_{\mathrm{dh}} = \left(\begin{array}{ccc|c} 0 & \ldots & 0 & 0 \\ \hline 1 & & 0 & 0 \\ & \ddots & & \vdots \\ 0 & & 1 & 0 \end{array}\right) \tag{19}$$

$$B_{\mathrm{dh}} = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

$$C_{\mathrm{dh}} = \begin{pmatrix} h[1] & \ldots & h[n] \end{pmatrix}$$

$$D_{\mathrm{dh}} = 0 \tag{20}$$

Note that the matrix A_{dh} has all its poles at the origin, as required for moving-average systems.

Eventually the model $M_{\mathrm{dh}} = (A_{\mathrm{dh}}, B_{\mathrm{dh}}, C_{\mathrm{dh}}, D_{\mathrm{dh}})$ has to be transformed into a low-order continuous-time model. At this stage, it may however not easily be converted into a continuous-time model, because of the zeros on the diagonal of the lower triangular matrix A_{dh}, which make it impossible to take a matrix logarithm. Therefore the model M_{dh} is first reduced in discrete-time and then converted to continuous-time, where at this point one has an accurate intermediate-order approximation. This intermediate-order model is then used to generate the (continuous-time) low-order starting points required for the optimization algorithm.

The steps 2-3 in Fig. 2 are similar to the procedure discussed in [5], where it is proposed to reduce the model M_{dh} by means of the popular *balance and truncate* procedure, first introduced by [8] and [9]. See also [10] for an interpretation of balancing in terms of energy passing from the input signal into the states and from the states onto the output signal.

The high-order discrete-time model M_{dh} is balanced first. A discrete-time state-space system (A,B,C,D) is balanced if the solutions P and Q of the discrete-time Lyapunov-Stein equations:

$$P - APA^T = BB^T \tag{21}$$

$$Q - A^T QA = C^T C \tag{22}$$

satisfy $P = Q = diag\{\sigma_1, \ldots, \sigma_n\}$ for some positive numbers σ_k, called the Hankel singular values of the system. Balanced realizations always exist and the Hankel singular values can

be ordered in a decreasing fashion. Observe that for the system M_{dh} the Lyapunov equation (21) only involves A_{dh} and B_{dh} and not the wavelet $\widetilde{\psi}(t)$ to which M_{dh} was fitted. It is not difficult to see that the corresponding solution P, called the *controllability Grammian*, is the identity matrix I_n.

Any state-space transformation T which transforms (A,B,C,D) into (TAT^{-1},TB,CT^{-1},D) takes the controllability Grammian P into TPT^T. Therefore, P remains equal to I_n if T is chosen to be an orthogonal matrix. The matrix Q, called the *observability Grammian*, is transformed by T into $T^{-T}QT^{-1}$. Since Q is symmetric and positive definite, it is possible to choose an orthogonal matrix T which makes the observability Grammian diagonal, e.g. by applying singular value decomposition to Q. It then is straightforward to compute an additional diagonal state-space transformation matrix which produces a balanced realization $(\bar{A}_{\mathrm{dh}},\bar{B}_{\mathrm{dh}},\bar{C}_{\mathrm{dh}},\bar{D}_{\mathrm{dh}})$. The procedure discussed in [5] arrives at the same balanced realization in a slightly different way, by observing that for the controllable companion form realization of a moving-average system, the associated Hankel matrix $\mathscr{H}$ built from the finite impulse response $0,h[1],h[2],\ldots,h[n]$ as

$$\mathscr{H} = \begin{pmatrix} h[1] & h[2] & h[3] & \ldots & h[n] \\ h[2] & h[3] & & h[n] & 0 \\ h[3] & & & 0 & \vdots \\ \vdots & h[n] & 0 & & \vdots \\ h[n] & 0 & \ldots & \ldots & 0 \end{pmatrix} \tag{23}$$

is diagonalized by the same matrix T. This avoids the computation of Q, of which the condition number is the square of the condition number of $\mathscr{H}$.

The balanced model (with its Hankel singular values in decreasing order) can now be reduced by truncating all the state variables corresponding to singular values σ_k below a certain threshold ε_σ, yielding a reduced (intermediate) order discrete-time system $M_{dm} = (A_{dm},B_{dm},C_{dm},0)$. The matrix A_{dm} will in general no longer have poles at the origin and therefore the system M_{dm} can be converted to a continuous-time system $M_{\mathrm{cm}} = (A_{\mathrm{cm}},B_{\mathrm{cm}},C_{\mathrm{cm}},0)$ by making use of the following relationships (see e.g. [11]):

$$A_{\mathrm{dm}} = e^{A_{\mathrm{cm}}\Delta t} \tag{24}$$

$$B_{\mathrm{dm}} = \int_0^{\Delta t} e^{A_{\mathrm{cm}}\tau} B_{\mathrm{cm}} d\tau \tag{25}$$

$$C_{\mathrm{dm}} = C_{\mathrm{cm}} \tag{26}$$

with A_{dm}, A_{cm}, B_{dm} and B_{cm} further conveniently related by (see [12]):

$$N = \begin{pmatrix} A_{\mathrm{cm}} & B_{\mathrm{cm}} \\ 0 & 0 \end{pmatrix} \tag{27}$$

$$e^{N\Delta t} = \begin{pmatrix} A_{\mathrm{dm}} & B_{\mathrm{dm}} \\ 0 & 1 \end{pmatrix} \tag{28}$$

The conversion from discrete-time to continuous-time starts from the right-hand side of Eq. (28) and involves taking the matrix logarithm.

The continuous-time model M_{cm} with intermediate order n_{m} must be designed to have an impulse response which closely matches the (time-reversed, time-shifted) wavelet function. It is subsequently used to generate low order approximations for various orders that can be used as starting points for the L_2-approximation procedure.

This can again be achieved with a balance and truncate procedure, but now in the continuous-time case, which involves the solution of continuous-time Lyapunov equations instead. The advantage of this layered set-up is that the expensive model reduction step from the high-order discrete-time moving-average approximation to the intermediate-order continuous-time approximation is carried out only once, while the cheaper model reduction to various low orders is carried out several times.

Any low-order continuous-time model M_{cr} obtained in this way, will in general not have a state-space form as displayed in (10) and it will not obey the constraint described in (15). To impose a step response which tends to zero, the constant term in the numerator of the transfer function of M_{cr} is set to zero. To bring the model into a form as displayed in (10), the complex eigenvalues of A_{cr} are ordered in complex pairs; the states corresponding to the real eigenvalues follow last. The new matrix A will have a block structure, from which the parameters required to solve the parametric optimization problem (11) can be easily read off. Note that A_{cr} can be used to fix the number of complex pairs of poles that is built into the parameterized model class.

V. RESULTS

In order to evaluate the new approach it was compared, with numerical simulations in Matlab, to the Padé-based approach in [1], where the Gaussian wavelet was approximated. Both the L_1 and the L_2-norm of the misfit are calculated (showing a lower value for a better fit), but only the L_2-norm was used as the optimization criterion. The performance is also compared to the truncated wavelet $\widetilde{\psi}(t)$, keeping in mind that this "wavelet" does not satisfy the requirement of a zero integral. The L_2-approximation was calculated for order 4 (L2O4) and order 5 (L2O5).

Furthermore the wavelet transform of an actual dataset was computed. The dataset used is a 24 hour, 128Hz sampled episode of a heart signal from PhysioNet's MIT-BIH nsrdb (normal sinus rhythm database); see [13]. The systems were simulated for scales $\sigma = 1,2,4,8,16,32$. The dataset used was lead 1 of episode 16539, which was split up into 15 batches, because of memory requirements, of which the results were averaged. The wavelet transformed signal $W_\psi(t,\sigma)$ at scales σ can be seen as a matrix with a row for each scale and a column for each time sample. By substracting the wavelet transform matrix that was obtained with the ideal wavelet from the matrix that was obtained by the approximating linear system, an error matrix was obtained. For this error matrix various norms can be computed, for which a lower value implies a better fit:

1) The matrix 1-norm or *column sum norm*, defined as the maximum of the 1-norms of the columns of a matrix

TABLE II

EVALUATION OF APPROXIMATIONS OF THE GAUSSIAN WAVELET

Criterion	Truncated	PadeO5	L2O4	L2O5
L_1 norm	0	0.2920	0.3675	0.1168
L_2 norm	0	0.1660	0.1476	0.0475
Col. sum norm	2590	675	618	560
Row sum norm	$8.43 \cdot 10^7$	$1.00 \cdot 10^7$	$9.06 \cdot 10^6$	$6.99 \cdot 10^6$
Spectral norm	$1.47 \cdot 10^5$	$3.43 \cdot 10^4$	$3.23 \cdot 10^4$	$2.99 \cdot 10^4$
Frobenius norm	$2.34 \cdot 10^5$	$4.33 \cdot 10^4$	$4.03 \cdot 10^4$	$3.41 \cdot 10^4$

TABLE III

L_2-ERRORS OF VARIOUS DAUBECHIES WAVELET APPROXIMATIONS

Order approx.	db2	db3	db4	db5	db6
8	0.1167	0.0965	0.1581	0.2141	0.2482
10	0.1051	0.0909	0.0848	0.1304	0.2987
12	0.1050	0.0812	0.0497	0.0755	0.1093
16	0.0877	0.0395	0.0387	0.0373	0.0334

2) The matrix ∞-norm or *row sum norm*, defined as the maximum of the 1-norms of the rows of a matrix
3) The matrix 2-norm or *spectral norm*, defined as the largest singular value of a matrix
4) The *Frobenius norm* or Euclidean norm, defined as the square root of the sum of squares of all the matrix entries

The results are listed in Table II, from which the superiority of the current approach is clearly visible.

The big practical advantage over the approach in [4] is that a starting point is now automatically generated, making it much easier to approximate wavelets. Previously a lot of tweaking of the starting point was required to obtain a satisfying approximation. With this approach also Daubechies wavelets have been approximated in [14]. Here the Daubechies wavelets were obtained from the filters with an iteration scheme (see for example [15]). The results of the approximation of various Daubechies wavelets are listed in Table III. The low-order Daubechies wavelets (db2, db3, etc.) are not very smooth (hard to fit), but have few oscillations (few complex pairs needed). The high-order Daubechies wavelets (db8, etc.) are smooth, but have more oscillations. These properties make these wavelets a challenge to approximate.

VI. CONCLUSION

In this paper a largely automated method to approximate wavelet functions for implementation in analog circuits is presented. The approach performs well in the sense that it leads to relatively low order approximations with an acceptable L_2-norm of the misfit, especially when compared to the wavelet approximations obtained with Padé-approximation. The final performance of an analog circuit will depend on more factors though, such as the quality of the components and the peculiarities of the intended practical application producing the signals to be analyzed. With this automated approach, wavelets that could not be approximated well with the approach in [1], such as the Daubechies wavelets, now were relatively easy to approximate well to a satisfactory degree.

REFERENCES

[1] S. Haddad, N. Verwaal, R. Houben, and W. Serdijn, "Optimized dynamic translinear implementation of the gaussian wavelet transform," in *Proc. IEEE International Symposium on Circuits and Systems (ISCAS'04)*, vol. I, May 2004, pp. 145–148.

[2] J. Mulder, W. Serdijn, A. C. van der Woerd, and A. H. van Roermund, "Dynamic translinear circuits - an overview," *Analog Integrated Circuits and Signal Processing*, vol. 22, no. 2-3, pp. 111–126, 2000.

[3] M. N. El-Gamal and G. W. Roberts, "A 1.2-V N-P-N-only integrator for log-domain filtering," *IEEE Trans. Circuits Syst. II*, vol. 49, no. 4, pp. 257–265, Apr. 2002.

[4] J. Karel, R. Peeters, R. Westra, S. Haddad, and W. Serdijn, "Wavelet approximation for implementation in dynamic translinear circuits," in *Proc. 16th IFAC World Congress (IFAC W.C.'05), Prague, Czech republic, July 3-8*, 2005.

[5] B. Beliczynski, I. Kale, and G. Cain, "Approximation of FIR by IIR digital filters: an algorithm based on balanced model reduction," *IEEE Trans. Signal Processing*, vol. 40, no. 3, pp. 532–542, Mar. 1992.

[6] T. Kailath, *Linear Systems*. Prentice Hall Inc., 1980.

[7] S. Mallat, *A Wavelet Tour of Signal Processing*. Academic Press, 1999.

[8] B. Moore, "Principal component analysis in linear systems: Controllability, observability, and model reduction," *IEEE Trans. Automat. Contr.*, vol. 26, no. 1, pp. 17–32, 1981.

[9] L. Pernebo and L. Silverman, "Model reduction via balanced state space representations," *IEEE Trans. Automat. Contr.*, vol. 27, no. 2, pp. 382–387, 1982.

[10] J. C. Willems and P. Rapisarda, "Balanced state representation from higher-order models," in *Int. Conf. on differential equations (Equadiff'03), Hasselt, Belgium*. WorldScientific 2005, 2003.

[11] R. J. Vaccaro, *Digital Control: A State-Space Approach*, ser. McGraw-Hill Series in Electrical an Computer Engineering. McGraw-Hill, 1995.

[12] C. van Loan, "Computing integrals involving the matrix exponential," *IEEE Trans. Automat. Contr.*, vol. 23, no. 3, pp. 395–404, 1978.

[13] A. Goldberger, L. Amaral, L. Glass, J. Hausdorff, P. Ivanov, R. Mark, J. Mietus, G. Moody, C.-K. Peng, and H. Stanley, "Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. 215–220, June 2000. [Online]. Available: http://circ.ahajournals.org/cgi/content/full/101/23/e215

[14] J. Karel, R. Peeters, R. Westra, S. Haddad, and W. Serdijn, "L2-norm based wavelet approximation for analog implementation," Maastricht University, Tech. Rep. M 04-05, Dec. 2004.

[15] G. Strang and T. Nguyen, *Wavelets and Filter Banks*. Wellesley-Cambridge Press, 1996.