

WAVELET APPROXIMATION FOR IMPLEMENTATION IN DYNAMIC TRANSLINEAR CIRCUITS

J.M.H. Karel* R.L.M. Peeters* R.L. Westra*
S.A.P. Haddad** W.A. Serdijn**

* *Universiteit Maastricht, Department of Mathematics,
P.O.Box 616, 6200 MD Maastricht, The Netherlands*

** *Technische Universiteit Delft, Faculty of Information
Technology and Systems, Electronics Research Laboratory,
Mekelweg 4, 2628 CD Delft, The Netherlands*

Abstract: For applications requiring low power consumption, signal processing in the analog domain is preferable. Approximate implementations of wavelet transforms in analog hardware can be achieved with dynamic translinear circuits. The quality of such implementations depends on the accuracy of the corresponding wavelet approximations. A design trade-off involves the approximation accuracy versus the complexity (model order) of the implemented filter. First we discuss the technique of Padé approximation for obtaining wavelet approximations. Then we present the technique of L_2 -approximation, which is conceptually more attractive but computationally more demanding. These techniques are compared by means of a worked example, involving Gaussian wavelet approximation and real measurements of an ECG signal. The L_2 -approximation approach is shown to exhibit superior performance. *Copyright ©2005 IFAC*

Keywords: analog signals, integrated circuit design, L_2 -approximation, Padé approximation, wavelets.

1. LOW-POWER IMPLEMENTATION OF WAVELET TRANSFORMS

In implantable medical devices such as pacemakers power consumption, especially for sensing circuits, is a critical issue because battery life is limited. In order to perform digital signal processing, an A/D (analog to digital) converter is required to transfer analog sensor information to the digital domain. Depending on the number of bits used, this is a heavy power consuming operation. For power consumption considerations it therefore

is preferable to perform as many computations as possible in the analog domain.

In previous work dynamic translinear (DTL) circuits have been used to implement low power signal processing operations. In (Haddad *et al.*, 2004) a DTL approach is described which aims to implement the Gaussian wavelet transform in the analog domain for the purpose of cardiac signal analysis. The performance of such an implementation depends largely on the accuracy of the approximations involved in this approach. From a technological point of view, the quality of the hardware components used in the manufacturing process may have a considerable impact on the performance of the IC, but we will not go into

¹ This paper is part of the BioSens project that is being funded by STW (Project no. DTC 6418).

such issues here. From a conceptual point of view, one of the critical steps concerns the approximation of the Laplace transform of the (time-reversed and shifted) Gaussian wavelet function by means of a strictly proper rational function of low order. For this purpose the classical technique of Padé approximation has been proposed. See (Baker Jr., 1975; Bultheel and Van Barel, 1986) for more details and a survey.

In the present paper we first discuss the advantages and disadvantages offered by the Padé approximation technique and some of its generalizations. We then present a second method for obtaining DTL-implementable approximations of wavelet transforms, based on L_2 -approximation. While the two approaches discussed in this paper can be used to approximate a variety of functions and transforms, the focus in this paper will be entirely on wavelet transforms.

Wavelet transforms usually cannot be implemented exactly in analog electronic hardware. The IC design of linear filters of finite order is quite well understood though. If a time signal $f(t)$ is passed through a linear system, then $f(t)$ is convolved with the impulse response $h(t)$ of that linear system, producing the output signal $\int_{-\infty}^{\infty} f(\tau)h(t-\tau)d\tau$. On the other hand, the continuous wavelet transform $W(t, \sigma)$ of $f(t)$ associated with a given mother wavelet $\psi(t)$ on a scale σ , is obtained as the integral $\frac{1}{\sqrt{\sigma}} \int_{-\infty}^{\infty} f(\tau)\psi\left(\frac{\tau-t}{\sigma}\right) d\tau$. Therefore, the analog computation of $W(t, \sigma)$ can be achieved through the implementation of a linear filter of which the impulse response satisfies $h(t) = \frac{1}{\sqrt{\sigma}}\psi(-t/\sigma)$. For obvious physical reasons only the hardware implementation of (strictly) causal stable filters is feasible. In other words, an implementable linear filter will have a (strictly) proper rational transfer function $H(s)$ that has all its poles in the complex left half plane.

Note that $h(t)$ will then be zero for negative t , so that any mother wavelet $\psi(t)$ which does not have this property must be time-shifted to facilitate an accurate approximation of its (correspondingly time-shifted) wavelet transform $W(t, \sigma)$. This may result in a truncation error for a wavelet that does not have compact support, such as the Gaussian wavelet. Note that an approximation error will also be due to the fact that a wavelet does not usually possess a rational Laplace transform.

Padé approximation provides a method to obtain a rational approximation, which for the intended application requires one to work in the frequency domain. Instead, the technique of L_2 -approximation can be formulated equally well in the time-domain and in the frequency domain, which offers some useful advantages as will be

demonstrated. The quality of the approximations resulting from these two approaches is evaluated on a practical level by comparing the resulting wavelet transformations on a data set of a heart signal. This shows that the DTL-implementable wavelet transform approximations obtained with the L_2 -approximation approach exhibit improved performance.

2. PADÉ APPROXIMATION OF WAVELET FUNCTIONS

In (Haddad *et al.*, 2004) Padé approximation is employed to approximate a given mother wavelet $\psi(t)$ in the Laplace domain by a suitable rational function $H(s)$. More precisely, $H(s)$ constitutes an approximation of the Laplace transform $\tilde{\Psi}(s)$ of the time-reversed and time-shifted mother wavelet $\tilde{\psi}(t) = \psi(t_0 - t)$, where t_0 denotes the time-shift involved.

Any Padé approximation $H(s)$ of $\tilde{\Psi}(s)$ is characterized by the property that the coefficients of the Taylor series expansion of $H(s)$ around a selected point $s = s_0$ coincide with the corresponding Taylor series coefficients of $\tilde{\Psi}(s)$ up to the highest possible order, given the pre-specified degrees of the numerator and denominator polynomials of $H(s)$. If we denote the Padé approximation $H(s)$ at $s = s_0$ and of order (m, n) with $m \leq n$ by

$$H(s) = \frac{p_0(s - s_0)^m + p_1(s - s_0)^{m-1} + \dots + p_m}{(s - s_0)^n + q_1(s - s_0)^{n-1} + \dots + q_n} \quad (1)$$

then there are $n + m + 1$ degrees of freedom, which generically makes it possible to match exactly the first $n + m + 1$ coefficients of the Taylor series expansion of $\tilde{\Psi}(s)$ around $s = s_0$. In fact, this matching problem can easily be rewritten as a system of $n + m + 1$ linear equations in the $n + m + 1$ variables $p_0, p_1, \dots, p_m, q_1, \dots, q_n$. See, e.g., (Baker Jr., 1975).

This brings us to one of the main advantages of Padé approximation: the linear system of equations will generically yield a unique solution which is easy to compute. Moreover, a good match is guaranteed between the given function $\tilde{\Psi}(s)$ and its approximation $H(s)$ in a neighborhood of the selected point s_0 . However, there are also some disadvantages which limit the practical applicability of this technique in the setting of this paper.

One important issue concerns the selection of the point s_0 . Note that a good approximation of $\tilde{\Psi}(s)$ over the entire (complex) Laplace domain is not a requirement *per se*. Instead, an approximation is needed which performs well when used for convolution in the time domain. Since the function $\tilde{\psi}(t)$ is a wavelet, it effectively will have compact




Fig. 1. Padé approximations of Gaussian wavelet

support and in particular it should be approximated well in the region of the time domain where it ‘lives’: which is somewhere near $t = 0$. Now, the initial value theorem for the Laplace transformation states that $\lim_{t \downarrow 0} h(t) = \lim_{s \rightarrow \infty} sH(s)$, which motivates the choice $s_0 = \infty$. This choice will lead to a good approximation of $\tilde{\psi}(t)$ near $t = 0$, as is demonstrated in Figure 1 for a 7th order approximation of the Gaussian wavelet.

A second important issue concerns stability. The approximation $h(t)$ of $\tilde{\psi}(t)$ is required to tend to zero for large values of t , since $\tilde{\psi}(t)$ will have this property too. However, stability does not automatically result from the Padé approximation technique. Indeed, if emphasis is put on obtaining a good fit near $t = 0$ by choosing $s_0 = \infty$, it may easily happen that the resulting approximation becomes unstable: see again the 7th order approximation in Figure 1. The selection of a suitable point s_0 involves a trade-off between a good fit near $t = 0$ and stability, yielding a non-trivial problem which may be difficult to handle, depending on the wavelet at hand. In this respect it may be of interest to note that the 5th order approximation described in (Haddad *et al.*, 2004) was obtained with the choice $s_0 = 0$ which actually corresponds to a good fit in the time domain for *large* values of t (see Figure 1).

A third issue concerns the choice of the degrees m and n of the numerator and denominator polynomials of the rational approximation $H(s)$. An unlucky choice may yield an inconsistent system of equations or an unstable approximation. Changing m or n may solve this problem, but the converse may also happen: one may run into stability problems even if s_0 is left unchanged.

In (Bultheel and Van Barel, 1986) an overview is given of various generalizations and extensions of Padé approximation which aim to deal with some of the problems just mentioned. For instance, it is possible to choose some or all of the poles of the rational approximation $H(s)$ in advance. This offers a possibility to deal with the stability issue, but no clear theory exists on the optimal choices for these poles. Another generalization involves the possibility to use more than one interpolation point, which for instance offers a method to deal

with the trade-off between $s_0 = \infty$ and $s_0 = 0$ in a more systematic way. Yet another possibility is to deal with many interpolation points s_0, s_1, \dots, s_k and to require only a match between the values of $H(s_i)$ and $\tilde{\Psi}(s_i)$ for $i = 0, 1, \dots, k$, no longer taking any derivatives at these points into account. One may even specify more interpolation points than the number of unknowns and use a linear least squares estimation technique to arrive at a unique solution. The advantage of such an approach is that one can optimize the function over a better distributed and controlled set of points than with classical Padé approximation. One may choose complex values here. A 9th order approximation of the Gaussian wavelet function obtained with this method is illustrated in Figure 1. This shows the feasibility of the approach, but for low order approximations the choice of interpolation points becomes more critical and the results are markedly less well.

However, all of these techniques remain to have one important drawback in common: the quality of the approximation of the wavelet is not measured directly in the time domain but in the Laplace domain, and the criterion used does not allow for a direct interpretation in system theoretic terms.

3. L_2 -APPROXIMATION OF WAVELET FUNCTIONS

The theory of L_2 -approximation provides an alternative framework for studying the problem of wavelet approximation which offers a number of advantages. On the conceptual level it is quite appropriate to use the L_2 -norm to measure the quality of an approximation $h(t)$ of the function $\tilde{\psi}(t)$. Indeed, the very definition of the wavelet transform itself involves the L_2 -inner product between the signal $f(t)$ and the mother wavelet $\psi(t)$. It is also desirable that the approximation $h(t)$ of $\tilde{\psi}(t)$ behaves equally well for all time instances t since $h(t)$ is used as a convolution kernel with any arbitrary shift. This property holds naturally for L_2 -approximation, but it is not supported by the Padé approximation approach.

Another advantage of L_2 -approximation is that it allows for a description in the time domain as well as in the Laplace domain, so that both frameworks can be exploited to develop further insight. According to Parseval’s equality the squared L_2 -norm of the difference between $\tilde{\psi}(t)$ and $h(t)$ can be expressed as:

$$\begin{aligned} \|\tilde{\psi} - h\|^2 &= \int_0^\infty (\psi(t) - h(t))^2 dt & (2) \\ &= \frac{1}{2\pi} \int_{-\infty}^\infty |\tilde{\Psi}(i\omega) - H(i\omega)|^2 d\omega. & (3) \end{aligned}$$

Minimization of $\|\tilde{\psi} - h\|$ is therefore equivalent to minimization of the L_2 -norm of the difference between the Laplace transforms $\tilde{\Psi}(s)$ and $H(s)$ over the imaginary axis $s = i\omega$. Note that this observation provides a rationale for the choice of interpolation points in a generalized Padé approximation approach.

One of the disadvantages of an L_2 -approximation approach is that there is a risk that the numerical optimization of $\|\tilde{\psi} - h\|$ ends in a local, non-global optimum. Several global optimization techniques and software packages exist but in general they provide no global optimality guarantee. Different starting points can give different local optima and thus can be used to find better solutions. Also, the outcomes of other approximation techniques can be used as starting points for L_2 -approximation.

Particularly in the case of low order approximation, the L_2 -approximation problem can be approached in a simple and straightforward way in the time domain. As is well known from linear systems theory (see, e.g., (Kailath, 1980)) 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:

$$\dot{x}(t) = Ax(t) + Bu(t), \quad (4)$$

$$y(t) = Cx(t), \quad (5)$$

where $u(t)$ denotes the scalar input to the filter, $y(t)$ the scalar filter output and $x(t) \in \mathbb{R}^n$ the state vector at time t . (Here, the direct feed-through is set to zero to achieve strict causality.) Then the impulse response function $h(t)$ and its Laplace transform $H(s)$ (i.e., the transfer function of the system) are given by:

$$h(t) = Ce^{At}B, \quad (6)$$

$$H(s) = C(sI - A)^{-1}B. \quad (7)$$

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. For low order systems, this makes it possible to propose an explicitly parameterized class of impulse response functions among which to search for a good approximation of $\psi(t)$. For instance, if a 5th order approximation is attempted, this parameterized class of functions $h(t)$ may typically have the following form:

$$h(t) = p_1 e^{p_6 t} + p_2 e^{p_7 t} \sin(p_8 t) + p_3 e^{p_7 t} \cos(p_8 t) + p_4 e^{p_9 t} \sin(p_{10} t) + p_5 e^{p_9 t} \cos(p_{10} t), \quad (8)$$

where the parameters p_6 , p_7 and p_9 must be strictly negative for reasons of stability. Note that

wavelets typically are oscillatory functions so that a good fit requires the contribution of sufficiently many damped harmonics, which further explains the structure of this class.

For the purpose of IC design it is useful to have an associated state-space representation (A, B, C) available. Such a representation is for instance provided by:

$$A = \begin{pmatrix} p_6 & 0 & 0 & 0 & 0 \\ 0 & p_7 & p_8 & 0 & 0 \\ 0 & -p_8 & p_7 & 0 & 0 \\ 0 & 0 & 0 & p_9 & p_{10} \\ 0 & 0 & 0 & -p_{10} & p_9 \end{pmatrix}, \quad B = \begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \\ 1 \end{pmatrix}, \quad (9)$$

$$C = (p_1 \ p_2 \ p_3 \ p_4 \ p_5).$$

Given the explicit form of the wavelet $\tilde{\psi}(t)$ and the parameterized class of functions $h(t)$, the L_2 -norm of the difference $\tilde{\psi} - h$ can now be minimized in a straightforward way using standard numerical optimization techniques and software. The negativity constraints on p_6 , p_7 and p_9 which enforce stability are not difficult to handle.

4. FURTHER WAVELET RELATED ISSUES

One common property of a wavelet function $\tilde{\psi}(t)$ that was undiscussed so far is that its integral is usually equal to zero: $\int_0^\infty \tilde{\psi}(t) dt = 0$. If this property is not shared by the approximation $h(t)$, this will cause an unwanted bias in the approximation of the wavelet transform. This is likely to happen in a situation where a truncation error occurs.

In terms of 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 . Indeed, if the wavelet transform is computed for a step input signal, then a bias will be manifest if such a property is not satisfied.

From the properties of the Laplace transform it holds that $\lim_{t \rightarrow \infty} \int_0^t h(\tau) d\tau = \lim_{s \downarrow 0} H(s)$. Therefore, the desired property comes down to $H(0) = 0$. In terms of a state-space representation (A, B, C) we have that

$$H(0) = -CA^{-1}B. \quad (10)$$

As an example, for the representation (9) it is not difficult to compute A^{-1} since it is block diagonal. This yields the explicit condition:

$$\frac{p_1}{p_6} + \frac{-p_2 p_8 + p_3 p_7}{p_7^2 + p_8^2} + \frac{-p_4 p_{10} + p_5 p_9}{p_9^2 + p_{10}^2} = 0. \quad (11)$$

If such an extra nonlinear condition is not conveniently handled by the optimization software, then it can easily be used to eliminate one of the variables from the problem.

5. EXAMPLE: L_2 -APPROXIMATION OF THE GAUSSIAN WAVELET

To demonstrate the L_2 -approximation technique we consider approximation of the Gaussian wavelet, which is a favorite choice in many biomedical wavelet applications; see, e.g., (Sahambi *et al.*, 1997). The outcomes are compared to the approximation obtained in (Haddad *et al.*, 2004) using Padé approximation. The Gaussian wavelet $\psi(t)$ is defined to be proportional to the first derivative of a Gaussian probability density function $\phi(t)$: $\phi(t) = (\frac{2}{\pi})^{\frac{1}{4}} e^{-t^2}$ and $\psi(t) = -2(\frac{2}{\pi})^{\frac{1}{4}} te^{-t^2}$. If one chooses the time-shift $t_0 = 2$, this gives rise to the following time-reversed and time-shifted wavelet function:

$$\tilde{\psi}(t) = -2 \left(\frac{2}{\pi}\right)^{\frac{1}{4}} (2-t)e^{-(2-t)^2}. \quad (12)$$

This function is truncated at $t = 0$, since only causal filters can be implemented. Note that the truncated wavelet $\tilde{\psi}(t)$ no longer has a zero integral.

The choice of the time shift $t_0 = 2$ involves an important trade-off that should be made with care. If t_0 is chosen too small, the truncation error will become too large. On the other hand, if t_0 is chosen too large, the function to be approximated will become very flat near $t = 0$. This effectively introduces a time-delay, which implies that a good fit can only be achieved with a filter of high order. However, the energy consumption of the IC increases with the order of the filter, so one of the prime design goals is to keep this order as small as possible.

To obtain a stable 5th order approximation of $\tilde{\psi}(t)$ the L_2 -approximation technique was applied using the parameterized class of functions given by Eqn. (8), where the stability constraints on p_6 , p_7 and p_9 and the additional constraint of Eqn. (11) were taken into account. The outcomes of the numerical optimization procedure depend on the choice of the starting point, for which several values were tried. The following filter was obtained:

$$H(s) = \frac{0.012s^4 - 1.567s^3 + 3.564s^2 - 30.1s}{s^5 + 4.069s^4 + 16.42s^3 + 30.0s^2 + 38.07s + 18.89}. \quad (13)$$

The corresponding wavelet approximation $h(t)$ is shown in Figure 2, yielding an L_2 -approximation error equal to 0.0472.

By fitting the function $\tilde{\psi}(t)$ with a 7th order model with three damped harmonics, using the same L_2 -approximation technique, an L_2 -approximation error equal to 0.00697 was obtained. This 7th order approximation is given by $H(s) =$

$$\frac{-0.061s^6 - 1.06s^5 - 3.93s^4 - 58.2s^3 - 14.4s^2 - 859s}{s^7 + 9.2s^6 + 57s^5 + 224s^4 + 599s^3 + 1065s^2 + 1132s + 566}. \quad (14)$$

Its wavelet approximation is shown in Figure 2.




Fig. 2. Approximations of the Gaussian wavelet.

6. PERFORMANCE EVALUATION

To determine the quality of the wavelet approximations obtained with the various techniques, one may resort to the computation of their L_2 -approximation errors. However, this will favor the L_2 -approximation approach, and more importantly it does not provide direct insight in the quality of the wavelet transforms that are obtained with the approximations $h(t)$ when applied to an actual data set.

Recall that $h(t)$ is used to approximate the wavelet function $\tilde{\psi}(t) = \psi(t_0 - t)$. For a given time signal $f(t)$, the approximate wavelet transform $W_h(t, \sigma)$ obtained with $h(t)$ at a certain scale σ is therefore given by:

$$W_h(t, \sigma) = \frac{1}{\sqrt{\sigma}} \int_{-\infty}^{\infty} f(\tau) h\left(t_0 + \frac{t - \tau}{\sigma}\right) d\tau. \quad (15)$$

To evaluate the similarity between the intended wavelet transforms $W(t, \sigma)$ and their approximations $W_h(t, \sigma)$, we may discretize the time axis and consider a number of different scales σ . In this way, $W(t, \sigma)$ and $W_h(t, \sigma)$ are stored as matrices of wavelet coefficients, which can conveniently be compared using matrix norms.

Four different matrix norms are considered in this paper: (i) the matrix 1-norm or *column sum norm*, defined as the maximum of the 1-norms of the columns of a matrix; (ii) the matrix ∞ -norm or *row sum norm*, defined as the maximum of the 1-norms of the rows of a matrix; (iii) the matrix 2-norm or *spectral norm*, defined as the largest singular value of a matrix; (iv) the *Frobenius norm* or Euclidean norm, defined as the square root of the sum of squares of all the matrix entries.

Another important issue for the evaluation of a wavelet transform approximation concerns the location of the modulus maxima. This aspect is tested by comparing the modulus maxima ridges in the original Gaussian wavelet transform with the ridges in the approximation, using a triangular weight function which discounts the distance between a reference ridge and its approximation.

The evaluation criteria discussed in this section have been calculated for three different wavelet

transform approximations: (a) the approximation (PadeO5), obtained in (Haddad *et al.*, 2004) using a 5th order Padé approximation. The approximation in that paper is performed on a non-time-reversed Gaussian wavelet, ignoring the constant $(\frac{2}{\pi})^{\frac{1}{4}}$ and setting the constant term in the Laplace transform of the wavelet function to zero. This gives the following transfer function:

$$H(s) = \frac{-(\frac{2}{\pi})^{\frac{1}{4}}(5.75s^3 - 18.3s^2 + 92.4s)}{s^5 + 8.33s^4 + 33.0s^3 + 74.8s^2 + 94.5s + 52.3}; \quad (16)$$

(b) the approximation (L2O5) obtained with the 5th order L_2 -approximation of Eqn. (13); (c) the approximation (L2O7) obtained with the 7th order L_2 -approximation of Eqn. (14).

As an actual data set to compute the wavelet transform we have used a 30 minutes, 360Hz sampled episode of a heart signal from PhysioNet's nstdb database; see (Goldberger *et al.*, 2000). The systems were simulated for scales $\sigma = 1, 2, 4, 8, 16, 32$. The results are listed below.

Table 1. Evaluation of approximations.

Criterion	PadeO5	L2O5	L2O7
L_2 norm	0.166	0.0472	0.00697
Col. sum	636.9	409.5	398.7
Row sum	$2.24 \cdot 10^7$	$8.51 \cdot 10^6$	$6.15 \cdot 10^6$
Spectral	$4.49 \cdot 10^4$	$2.20 \cdot 10^4$	$2.10 \cdot 10^4$
Frobenius	$5.25 \cdot 10^4$	$2.78 \cdot 10^4$	$2.51 \cdot 10^4$
Mod. max.	$1.24 \cdot 10^6$	$1.27 \cdot 10^6$	$1.31 \cdot 10^6$

For all but the modulus maxima measure a lower value indicates a better performance. This shows that L_2 -approximation gives superior results over Padé approximation for this application.

7. CONCLUSIONS

As we have seen, the Padé approximation approach has the advantages that it is a fast algorithm and that it yields a unique solution. However the selection of the orders m and n is not a straightforward task. Furthermore the selection of the expansion point s_0 is a trade-off between stability and a good fit around zero in the time domain. Extensions and generalizations of this technique can partially overcome these problems, but they all have the drawback that the quality of the approximation of the wavelet is measured in the Laplace domain instead of in the time domain and it has no direct physical interpretation. Increasing the order of the approximation will make it easier to find a good approximation; however this will result in an increase in power consumption that may not be acceptable for the intended application.

The proposed L_2 -approximation technique is appropriate since wavelet transforms are L_2 -inner

products, all time instances are treated equally well and it has a convenient reformulation in the Laplace domain. The fit is performed directly in the time domain yielding good control and easy interpretation of the optimization criteria. With this approach it is possible to approximate wavelets that can not be approximated sufficiently well with an acceptable order in a straightforward way with the Padé approach. The drawbacks of this technique are that it is computationally more demanding than the Padé approach and that no guarantees for global optimization exist.

A possible approach to select a starting point is to first compute a high-order close approximation (which should be easy to obtain) and then to perform model reduction, for instance with the popular balance and truncate procedure. This approach is currently under investigation.

REFERENCES

- Baker Jr., G.A. (1975). *Essentials of Padé Approximants*. Academic Press.
- Bultheel, A. and M. Van Barel (1986). Padé techniques for model reduction in linear system theory: a survey. *Journal of Computational and Applied Mathematics* **14**, 401–438.
- Daubechies, I. (1992). *Ten Lectures on Wavelets*. SIAM.
- Goldberger, A.L., L.A.N. Amaral, L. Glass, J.M. Hausdorff, P.Ch. Ivanov, R.G. Mark, J.E. Mietus, G.B. Moody, C.-K. Peng and H.E. Stanley (2000). Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals. *Circulation* **101**(23), e215–e220.
- Haddad, S.A.P., N. Verwaal, R. Houben and W.A. Serdijn (2004). Optimized dynamic translinear implementation of the gaussian wavelet transform. In: *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS)*. Vol. I. pp. 145–148.
- Kailath, T. (1980). *Linear Systems*. Prentice Hall.
- Leblond, J. and M. Olivi (1998). Weighted H2-approximation of transfer functions. *Mathematics of Control, Signals and Systems* **11**, 28–39.
- Mallat, S. (1999). *A Wavelet Tour of Signal Processing*. Academic Press.
- Peeters, R.L.M., B. Hanzon and D. Jibeteau (2003). Optimal H2-model reduction in state-space: a case study. In: *Proceedings of the European Control Conference 2003 (ECC03)*.
- Sahambi, J., S. Tandon and R. Bhatt (1997). Using wavelet transforms for ECG characterization. *IEEE Engineering in Medicine and Biology Magazine* **16**, 77–83.
- Strang, G. and T. Nguyen (1996). *Wavelets and Filter Banks*. Wellesley-Cambridge Press.