

Efficient Compression of QRS Complexes Using Hermite Expansion

Aliaksei Sandryhaila, *Member, IEEE*, Samir Saba, Markus Püschel, *Senior Member, IEEE*, and Jelena Kovacevic, *Fellow, IEEE*

Abstract—We propose a novel algorithm for the compression of ECG signals, in particular QRS complexes. The algorithm is based on the expansion of signals with compact support into a basis of discrete Hermite functions. These functions can be constructed by sampling continuous Hermite functions at specific sampling points. They form an orthogonal basis in the underlying signal space. The proposed algorithm relies on the theory of signal models based on orthogonal polynomials. We demonstrate that the constructed discrete Hermite functions have important advantages compared to continuous Hermite functions, which have previously been suggested for the compression of QRS complexes. Our algorithm achieves higher compression ratios compared with previously reported algorithms based on continuous Hermite functions, discrete Fourier, cosine, or wavelet transforms.

Index Terms—Compression, ECG signal, Hermite function, Hermite transform, orthogonal polynomials, QRS complex, signal model.

I. INTRODUCTION

SOME classes of electrophysiological signals have (or can be assumed to have) compact support. These signals represent the impulse responses of a system or an organ to an electrical stimulation that is recorded on the body surface. Examples include electrocardiographic (ECG), electroencephalographic, and myoelectric signals.

The major role of electrophysiological signals is to provide information about a patient's condition, and reflect important changes in his/her state. In addition, it may also be desired to store electrophysiological signals for later analysis and reference. However, the visual analysis and monitoring of long-term repetitive signals is a tedious task that requires the presence of a human operator. In these cases, computer-based systems can be used to facilitate this process.

Manuscript received February 16, 2011; revised July 14, 2011 and September 30, 2011; accepted October 13, 2011. Date of publication October 25, 2011; date of current version January 13, 2012. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Yufei Huang. This work was supported in part by NSF grant CCF-0634967.

A. Sandryhaila is with the Department of Electrical and Computer Engineering, Carnegie-Mellon University, Pittsburgh, PA 15213 USA (e-mail: asandryh@andrew.cmu.edu).

S. Saba is with the Cardiovascular Institute, University of Pittsburgh Medical Center, Pittsburgh, PA 15213 USA (e-mail: sabas@upmc.edu).

M. Püschel is with the Department of Computer Science, ETH Zurich, Switzerland (e-mail: pueschel@inf.ethz.ch).

J. Kovacevic is with the Departments of Biomedical Engineering and Electrical and Computer Engineering, Carnegie-Mellon University, Pittsburgh, PA 15213 USA (e-mail: jelenak@cmu.edu).

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/TSP.2011.2173336

For efficient storage, automatic monitoring, and accurate interpretation, electrophysiological signals are usually represented by a set of features, either heuristic, such as duration and amplitude, or formal, such as the coefficients of the expansion in an orthogonal basis. In the latter case, one can use either a continuous or a discrete basis. When a continuous basis is used, the projection and reconstruction of a compactly supported signal are computed using numerical methods for integral approximation, for example, a numerical quadrature. When a discrete basis is used, a discrete signal transform, such as the discrete Fourier transform or the discrete cosine transform, is applied to a digitized signal. These signals can be obtained from continuous ones by sampling at specific sampling points. The choice of the quadrature formula, or the sampling points is an important characteristic of the representation scheme.

In both cases of continuous and discrete bases, typically only a few projection coefficients are used for the storage and reconstruction of a signal. While the reduction of expansion coefficients leads to the compression of the signal, it also leads to a reconstruction error. The goal of the compression optimization then is to minimize the error while maximizing the compression ratio. This can be achieved, for example, by using only the coefficients with the largest magnitude, and minimizing the number of coefficients used for reconstruction.

In this paper, we introduce a novel algorithm for the compression of *QRS complexes* (also known as QRS intervals), which are the most characteristic waves of ECG signals. The structure of an ECG signal is shown in Fig. 1. The morphology of QRS complexes is important to cardiologists in different stages of diagnosis and treatment [1]. Examples include, but are not limited to the following.

- 1) The detection of the rhythm origin (supraventricular, from the upper chambers of the heart, or ventricular, from the bottom chambers). This is a key factor in the choice of the treatment.
- 2) The detection of conduction abnormalities between atria and ventricles (bundle branch blocks, fascicular blocks, and others).
- 3) The detection of accessory pathways (Wolf-Parkinson-White syndrome).
- 4) The evaluation of the effect of specific antiarrhythmic medications (such as sodium channel blockers).

Our proposed algorithm is based on the expansion of QRS complexes into the basis of *discrete* Hermite functions. Such functions are obtained by sampling *continuous* Hermite functions at specific sampling points, not necessarily located on a uniform grid. We originally proposed this compression method

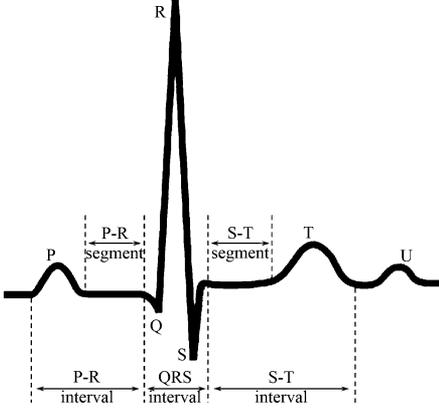


Fig. 1. The schematic structure of an ECG signal.

in [2]. In this paper, we extend our previous work by rigorously formulating and describing the signal model used for the description of QRS complexes. In particular, we use results from our recently developed theory of signal models based on orthogonal polynomials [3]–[5]. We also identify a fast algorithm for the implementation of the proposed compression method. Finally, our experimental results demonstrate that the proposed method achieves higher compression ratios, compared to other methods, when used to obtain medically acceptable compressed ECG signals. Here, we call a compressed ECG signal *medically acceptable*, if visually it is sufficiently similar to the original signal, and would not lead to an incorrect interpretation and diagnosis. In our experiments, the identification of medically acceptable compressed signals has been performed by one of the authors who is a cardiologist with extensive experience in cardiac electrophysiology.

Related work. Previous work on the compression of QRS complexes includes the use of continuous Hermite functions as the expansion basis [6]–[9]. Due to the shape similarity between continuous Hermite functions and QRS complexes, these functions were identified as a suitable basis for the representation and compression of the latter. These works, however, only provide a theoretical framework for the compression algorithm. The actual implementation is not discussed and no experimental data is provided that would confirm that the proposed compression method indeed performs better than other methods. In addition, as we discuss in Section II, the computer implementations of the previously proposed algorithms suffer from certain limitations, such as the inability to obtain an exact reconstruction of a signal, large computational cost, and an a priori selection of coefficients for reconstruction.

II. BACKGROUND

In this section, we discuss the expansion of continuous signals using Hermite functions, its digital implementation, and its use in signal compression.

A. Continuous Hermite Functions

Consider the family of polynomials $H_\ell(t)$, $\ell \geq 0$, that satisfy the recursion

$$H_\ell(t) = 2tH_{\ell-1}(t) - 2(\ell - 1)H_{\ell-2}(t) \quad (1)$$

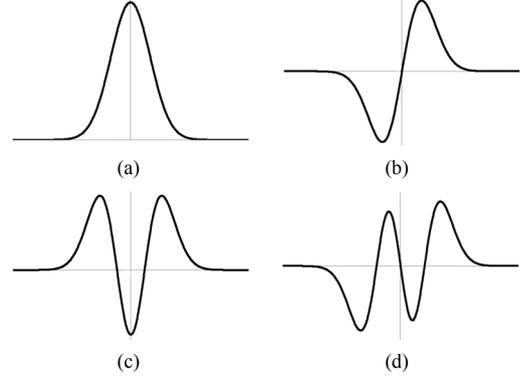


Fig. 2. First four Hermite functions (plotted for the same scale σ). (a) $\varphi_0(t, \sigma)$. (b) $\varphi_1(t, \sigma)$. (c) $\varphi_2(t, \sigma)$. (d) $\varphi_3(t, \sigma)$.

for $\ell \geq 2$, with $H_0(t) = 1$ and $H_1(t) = 2t$. They are known as *Hermite polynomials*. These polynomials are orthogonal on the real line \mathbb{R} with respect to the weight function e^{-t^2}

$$\int_{-\infty}^{\infty} H_\ell(t)H_m(t)e^{-t^2} dt = 2^\ell \ell! \sqrt{\pi} \cdot \delta_{\ell-m}. \quad (2)$$

It immediately follows from (2) that the functions

$$\varphi_\ell(t, \sigma) = \frac{1}{\sqrt{\sigma 2^\ell \ell! \sqrt{\pi}}} e^{-\frac{t^2}{2\sigma^2}} H_\ell\left(\frac{t}{\sigma}\right) \quad (3)$$

are orthonormal on \mathbb{R} with respect to the standard inner product

$$\langle \varphi_\ell(t, \sigma), \varphi_m(t, \sigma) \rangle = \int_{\mathbb{R}} \varphi_\ell(t, \sigma) \varphi_m(t, \sigma) dt = \delta_{\ell-m}. \quad (4)$$

The set of functions $\{\varphi_\ell(t, \sigma)\}_{\ell \geq 0}$, called *continuous Hermite functions*, is an orthonormal basis in the Hilbert space of continuous functions defined on \mathbb{R} [10]–[13]. Any such function $s(t)$ can be represented as a linear combination of the basis functions

$$s(t) = \sum_{\ell \geq 0} c_\ell \varphi_\ell(t, \sigma) \quad (5)$$

where

$$c_\ell = \langle s(t), \varphi_\ell(t, \sigma) \rangle = \int_{\mathbb{R}} s(t) \varphi_\ell(t, \sigma) dt.$$

The first four continuous Hermite functions are shown in Fig. 2. Note that each $\varphi_\ell(t, \sigma)$ quickly approaches zero as the value of $|t|$ increases: since $H_\ell(\frac{t}{\sigma})$ is a polynomial of degree ℓ

$$\lim_{|t| \rightarrow \infty} e^{-\frac{t^2}{2\sigma^2}} H_\ell\left(\frac{t}{\sigma}\right) = 0.$$

As a consequence, for practical purposes we can assume that each continuous Hermite function has a compact support. Since in this paper we often work only with the first L Hermite functions, we assume that $\varphi_0(t, \sigma), \varphi_1(t, \sigma), \dots, \varphi_{L-1}(t, \sigma)$ have the same compact support $[-T_\sigma, T_\sigma]$, where T_σ is a suitably chosen constant that depends on σ and L . In other words, we assume

$$\varphi_\ell(t, \sigma) = 0 \text{ for } t \notin [-T_\sigma, T_\sigma]$$

where $0 \leq \ell < L$. If a signal $s(t)$ also has compact support $[-T_\sigma, T_\sigma]$, then we can compute the coefficients c_ℓ using a finite integral

$$c_\ell = \int_{\mathbb{R}} s(t) \varphi_\ell(t, \sigma) dt = \int_{-T_\sigma}^{T_\sigma} s(t) \varphi_\ell(t, \sigma) dt. \quad (6)$$

B. Compression With Continuous Hermite Functions

Coefficient-based compression. In practical applications, only a finite number M of Hermite functions are used to represent the signal $s(t)$ in (5). Accordingly, only a few coefficients $c_{\ell_0}, \dots, c_{\ell_{M-1}}$ need to be computed. Here, c_{ℓ_m} corresponds to $\varphi_{\ell_m}(t, \sigma)$ in (6). The approximated signal is then

$$\hat{s}(t) = \sum_{m=0}^{M-1} c_{\ell_m} \varphi_{\ell_m}(t, \sigma). \quad (7)$$

Alternatively, a larger pool of coefficients can be computed, from which M optimal ones are selected. It is well-known that, in an orthonormal basis, selecting coefficients with the largest magnitude minimizes the approximation error computed as the energy of the difference between the signal $s(t)$ and its approximation with M basis functions. If a basis is orthogonal, but not orthonormal (i.e., basis functions do not have unit norms), an additional weighting of coefficients proportionally to the norms of the basis functions can be performed.

Digital implementation. The coefficients c_ℓ in (6) and the Hermite expansion (5) are computed using the continuous functions. However, in practice they have to be computed in a digital form.

In particular, for each coefficient c_{ℓ_m} the integral in (6) can be calculated using a numerical quadrature that is based, for example, on a rectangle rule

$$\begin{aligned} c_{\ell_m} &= \int_{-T_\sigma}^{T_\sigma} s(t) \varphi_{\ell_m}(t, \sigma) dt \\ &\approx \sum_{k=-K}^K s(\tau_k) \varphi_{\ell_m}(\tau_k, \sigma) (t_k - t_{k-1}). \end{aligned} \quad (8)$$

Here, $-T = t_{-K-1} < t_{-K} < \dots < t_{K-1} < t_K = T$. Each sampling point τ_k is located inside the corresponding interval: $t_{k-1} \leq \tau_k \leq t_k$.

Then, we can compute the discrete version of the approximation signal $\hat{s}(t)$ it corresponds to $\hat{s}(t)$ sampled at points τ_k , $-K \leq k \leq K$

$$\hat{s}(\tau_k) = \sum_{m=0}^{M-1} c_{\ell_m} \varphi_{\ell_m}(\tau_k, \sigma). \quad (9)$$

Usually, points t_k are assumed to lie on a uniform grid, such that $t_k - t_{k-1} = \Delta$ for all k . Then (8) and (9) can be expressed in matrix-vector notation. Let us define vectors

$$\mathbf{s} = \begin{pmatrix} s(\tau_{-K}) \\ \vdots \\ s(\tau_K) \end{pmatrix}, \quad \mathbf{c} = \begin{pmatrix} c_0 \\ \vdots \\ c_{M-1} \end{pmatrix}, \quad \hat{\mathbf{s}} = \begin{pmatrix} \hat{s}(\tau_{-K}) \\ \vdots \\ \hat{s}(\tau_K) \end{pmatrix}. \quad (10)$$

Then

$$\mathbf{c} = \Delta \Phi^T \mathbf{s}, \quad (11)$$

$$\hat{\mathbf{s}} = \Phi \mathbf{c} \quad (12)$$

where $\Phi \in \mathbb{R}^{(2K+1) \times M}$, such that its m th column is the ℓ_m th Hermite function sampled at the points $\tau_{-K}, \tau_{-K+1}, \dots, \tau_K$

$$\Phi = \left[\varphi_{\ell_m}(\tau_{k-K}, \sigma) \right]_{0 \leq k < 2K+1, 0 \leq m < M}. \quad (13)$$

As follows from (11)–(12), to achieve the perfect reconstruction $\hat{\mathbf{s}} = \mathbf{s}$, Φ must satisfy $\Phi \Phi^T = I_{2K+1}$.

Compression of QRS complexes: Previous work. The compression of QRS complexes using the expansion into continuous Hermite functions has been studied in [6]–[9]. It was originally motivated by the visual similarity of QRS complexes, centered around their peaks, and Hermite functions, as can be observed from Figs. 1 and 2. Varying the value of σ corresponds to “stretching” or “compressing” Hermite functions $\varphi_\ell(t, \sigma)$ to optimally match a given QRS complex.

Since ECG signals are usually available as discrete signals equidistantly sampled at $\tau_k = k\Delta$, previously reported work assumed, explicitly or implicitly, $t_k = \tau_k = k\Delta$ in (8)–(9) and hence led exactly to the matrix-vector products in (11) and (12). In addition, it proposed to use only the *first* M Hermite functions $\varphi_0(t, \sigma), \dots, \varphi_{M-1}(t, \sigma)$ for the approximation of QRS complexes.

This compression algorithm has several important limitations. First, since $\Phi \Phi^T \neq I_{2K+1}$ for $\tau_k = k\Delta$, the approximation $\hat{\mathbf{s}}$ does not converge to the original signal \mathbf{s} , regardless of the number M of Hermite functions used for the construction of an approximation. As a result, the original signal \mathbf{s} cannot be reconstructed exactly.

This problem could be theoretically addressed by using $M = 2K + 1$ coefficients and replacing Φ^T with Φ^{-1} to compute coefficients \mathbf{c} in (11). However, the computation of Φ^{-1} is a nontrivial task. Moreover, the matrix-vector product $\Phi^{-1} \mathbf{s}$ requires $(2K + 1)^2$ operations. This cost can become prohibitive for large value of K and make this approach impractical.

Finally, the approximation of $\hat{s}(t_k)$ with the *first* $M < 2K + 1$ coefficients c_0, c_1, \dots, c_{M-1} in (8) may not be the optimal choice for the construction of approximation $\hat{\mathbf{s}}$ with M basis functions.

In Section IV, we propose an improved compression algorithm that samples ECG signals at nonequidistant points, and selects coefficients c_{ℓ_m} with the *largest* magnitude among a larger selection of coefficients.

III. SIGNAL MODEL FOR QRS COMPLEXES

In this section, we construct a signal model for the description of QRS complexes. It is based on the scaled Hermite polynomials. We identify the Fourier transform for this model, called a *Hermite transform*, as well as its inverse. We also describe a fast algorithm for the corresponding Hermite transform. These results will be used in Section IV to construct and implement the novel compression algorithm.

A. Signal Model

Let us define polynomials

$$P_\ell(t) = \frac{1}{\sqrt{2^\ell \ell!}} H_\ell(t), \quad (14)$$

where $H_\ell(t)$ denotes the ℓ th Hermite polynomials (1). Throughout the paper, we refer to polynomials $P_\ell(t)$, $\ell \geq 0$, as *scaled Hermite polynomials*.

Signal model. Consider the vector space \mathcal{M} of functions spanned by the first N scaled Hermite polynomials (14)

$$\mathcal{M} = \{s_0 P_0(t) + s_1 P_1(t) + \dots + s_{N-1} P_{N-1}(t)\} \quad (15)$$

where $(s_0, s_1, \dots, s_{N-1})^T \in \mathbb{R}^N$. Since each $P_\ell(t)$ is a polynomial of degree ℓ , \mathcal{M} is a vector space of degree N . It is closed under the addition and linear scaling of its elements.

Assume that \mathcal{M} is also closed under the multiplication of its elements modulo polynomial $P_N(t)$.¹ As we demonstrate in [3]–[5], in this case \mathcal{M} has exactly N spectral components, and the k th component of a signal

$$s(t) = s_0 P_0(t) + s_1 P_1(t) + \dots + s_{N-1} P_{N-1}(t)$$

is defined as

$$s(\alpha_k) = s_0 P_0(\alpha_k) + s_1 P_1(\alpha_k) + \dots + s_{N-1} P_{N-1}(\alpha_k).$$

where $\alpha_0, \alpha_1, \dots, \alpha_{N-1}$ are the roots of $P_N(t)$. Note that $\alpha_0, \alpha_1, \dots, \alpha_{N-1}$ are all distinct real numbers [10]. Without loss of generality, we assume $\alpha_0 < \alpha_1 < \dots < \alpha_{N-1}$.

Hermite transform. The spectral decomposition of signal $s(t)$ is defined as

$$S(\alpha) = (s(\alpha_0) \quad s(\alpha_1) \quad \dots \quad s(\alpha_{N-1}))^T. \quad (16)$$

We can express the spectral decomposition of $s(t)$ in the following matrix-vector notation. Consider the $N \times N$ matrix

$$\mathcal{P}_{P,\alpha} = [P_\ell(\alpha_k)]_{0 \leq k, \ell < N}. \quad (17)$$

The ℓ th column of $\mathcal{P}_{P,\alpha}$ is polynomials $P_\ell(t)$ evaluated at $\alpha_0, \dots, \alpha_{N-1}$. In general, matrices with this structure are called *polynomial transforms*. In the particular case of (17), We call $\mathcal{P}_{P,\alpha}$ the (*forward*) *Hermite transform*.²

¹Strictly speaking, we must consider two separate spaces \mathcal{A} and \mathcal{M} that have the same elements: $\mathcal{A} = \mathcal{M}$. \mathcal{A} is an algebra: it is a vector space closed under the addition and multiplication of its elements modulo $P_N(t)$. Then \mathcal{M} can be viewed as an \mathcal{A} -module, and its spectral decomposition can be defined.

\mathcal{A} and \mathcal{M} comprise a part of what we call an *algebraic signal model*. It is a central component of the *algebraic signal processing theory* developed by the authors in [3]–[5], [14]–[18].

In particular, in [3]–[5] we introduce signal models based on orthogonal polynomials (of which Hermite polynomials are a special case). A family of polynomials $\{P_\ell(t)\}_{\ell \geq 0}$ is called *orthogonal*, if they satisfy a recursion of the form $tP_\ell(t) = a_\ell P_{\ell-1}(t) + b_\ell P_\ell(t) + c_\ell P_{\ell+1}(t)$, usually with initial conditions $P_0(t) = 1$ and $P_{-1}(t) = 0$. Each family is orthogonal over an interval $I \subseteq \mathbb{R}$ with a weight function $w(t) : \int_I P_\ell(t) P_m(t) w(t) dt = 0$ if $\ell \neq m$. Each polynomial $P_\ell(t)$ has exactly ℓ simple real roots. A detailed discussion on orthogonal polynomials can be found in [10]–[13].

²Another definition of the Hermite transform has been previously used in [19], [20]. However, it denotes an expansion into continuous Hermite polynomials, and should not be confused with (17).

The spectral decomposition (16) can then be computed as the matrix-vector product

$$\begin{pmatrix} s(\alpha_0) \\ s(\alpha_1) \\ \vdots \\ s(\alpha_{N-1}) \end{pmatrix} = \mathcal{P}_{P,\alpha} \cdot \begin{pmatrix} s_0 \\ s_1 \\ \vdots \\ s_{N-1} \end{pmatrix}. \quad (18)$$

Inverse Hermite transform. In general, it is non-trivial to compute the inverse of the matrix $\mathcal{P}_{P,\alpha}$ for arbitrary polynomials $P_\ell(t)$. However, as follows from Theorem 16 in [3] and explicitly derived in [5], we can apply the Christoffel-Darboux formula for orthogonal polynomials [10] to the scaled Hermite polynomials $P_\ell(t)$ to obtain the inverse of the Hermite transform (17)³

$$\mathcal{P}_{P,\alpha}^{-1} = \mathcal{P}_{P,\alpha}^T \cdot \mathbf{D}. \quad (19)$$

Here, $\mathbf{D} \in \mathbb{R}^{N \times N}$ is a diagonal matrix with the k th diagonal element equal to

$$D_{k,k} = \frac{\sqrt{\frac{2}{N}}}{P_{N-1}'(\alpha_k) P_N'(\alpha_k)}$$

where $P_N'(t)$ denotes the derivative of $P_N(t)$.

B. Fast Algorithm for Hermite Transform

A straightforward computation of the matrix-vector product in (18) requires, in general, $2N^2$ additions and multiplications. Alternatively, $43N \log_2^2 N$ additions and multiplications can be required, if one uses an algorithm proposed in [22]. In both cases, the computational cost can become unacceptable for large values of N , especially if the product has to be computed in real-time (for example, in the case of ECG signal processing). The same applies to the inverse $\mathcal{P}_{P,\alpha}^{-1}$.

Here, we provide an improved computational algorithm for the matrix-vector product in (18). It reduces the computational cost approximately by a factor of two. We only state the algorithm here, without the proofs, since they can be verified by direct computation. We originally derived this algorithm by applying the theory of fast polynomial transforms developed in [4] and [23] to the signal space (15).

Theorem 1: Let $N = 2M$ be an even number. Then the polynomial transform (17) can be factored as

$$\mathcal{P}_{P,\alpha} = \begin{pmatrix} \mathbf{J}_M & -\mathbf{J}_M \\ \mathbf{I}_M & \mathbf{I}_M \end{pmatrix} \begin{pmatrix} \mathbf{I}_M & & & \\ & \alpha_M & & \\ & & \ddots & \\ & & & \alpha_{N-1} \end{pmatrix} \times \begin{pmatrix} \mathcal{P}_M & \\ & \mathcal{P}_M \end{pmatrix} \mathbf{B}_0 \mathbf{B}_1 \dots \mathbf{B}_{M-1} \mathbf{L}_2^N. \quad (20)$$

³In [21], weighted Hermite polynomials sampled at the roots of $P_N(t)$ have also been identified as an orthogonal basis, essentially leading to (19).

Here, \mathbf{I}_M is an identity matrix of size M and \mathbf{L}_M is a complementary identity matrix of size M

$$\mathbf{J}_M = \begin{pmatrix} & & & 1 \\ & & \ddots & \\ & 1 & & \\ 1 & & & \end{pmatrix}.$$

Matrix \mathcal{P}_M is an $M \times M$ polynomial transform defined as

$$\mathcal{P}_M = [P_{2\ell}(\alpha_{k+M})]_{0 \leq k, \ell < M}.$$

Each B_k is an identity matrix except its $(M+k, M+k)$ th and $(M+k-1, M+k)$ th elements are equal to $\sqrt{\frac{2}{(2k+1)}}$ and $-\sqrt{\frac{2k}{(2k+1)}}$, respectively. \mathbf{L}_2^N is a permutation matrix: its (k, ℓ) th element is 1 if $\ell = \lfloor \frac{2k(N+1)}{N} \rfloor \bmod N$ and 0 otherwise.

The algorithm for odd values of N is similar.

Theorem 2: Let $N = 2M + 1$ be an odd number. Then the polynomial transform (17) can be factored as

$$\mathcal{P}_{P,\alpha} = \begin{pmatrix} \mathbf{J}_M & -\mathbf{J}_M \\ \mathbf{I}_M & \mathbf{I}_M \end{pmatrix} \begin{pmatrix} \mathbf{I}_{M+1} & & & \\ & \alpha_{M+1} & & \\ & & \ddots & \\ & & & \alpha_{2M} \end{pmatrix} \times \begin{pmatrix} \mathcal{P}_{M+1} & \\ & \mathcal{P}_M \end{pmatrix} B_0 B_1 \dots B_{M-1} \mathbf{L}_2^N. \quad (21)$$

Here, \mathcal{P}_{M+1} is an $(M+1) \times (M+1)$ polynomial transform defined as

$$\mathcal{P}_{M+1} = [P_{2\ell}(\alpha_{k+M})]_{0 \leq k, \ell < M+1}.$$

\mathcal{P}_M is a submatrix of \mathcal{P}_{M+1} obtained by removing the first row and last column of the latter. Each B_k is an identity matrix except its $(M+1+k, M+1+k)$ th and $(M+k, M+1+k)$ th elements are equal to $\sqrt{\frac{2}{(2k+1)}}$ and $-\sqrt{\frac{2k}{(2k+1)}}$, respectively. Matrices \mathbf{I}_M , \mathbf{J}_M and \mathbf{L}_2^N are as described above.

We show in [4] that the polynomial transforms \mathcal{P}_M and \mathcal{P}_{M+1} in Theorems 1 and 2 are also based on orthogonal polynomials. Hence, they can also be computed using approximately $2M^2$ or $43M \log_2^2 M$ additions and multiplications, depending on the choice of the algorithm. Since other matrices in factorizations (20) and (21) require approximately $3N$ additions and multiplication, the computational algorithms in Theorems 1 and 2 require approximately two times fewer operations compared to other algorithms.

IV. COMPRESSION ALGORITHM

As we mentioned in Section II, the parameter σ in the definition of the continuous Hermite functions (3) can be used to “stretch” and “compress” the functions relatively to the signal $s(t)$.

Algorithm modifications. Alternatively, we can fix $\sigma = 1$, and introduce a parameter λ to “stretch” and “compress” signal

$s(t\lambda)$ instead. In this case the numerical quadrature (8) can be rewritten as

$$c_\ell = \int_{-T_\lambda}^{T_\lambda} s(t\lambda) \varphi_\ell(t, 1) dt \approx \sum_{k=-K}^K s(\tau_k \lambda) \varphi_\ell(\tau_k, 1) (t_k - t_{k-1}). \quad (22)$$

Furthermore, instead of sampling points on a uniform grid, we propose to use nonequispaced sampling points associated with the roots of Hermite polynomials. Namely, we set $\tau_k = \alpha_{k+K}$, $-K \leq k \leq K$, to be the roots of the scaled Hermite polynomial $P_{2K+1}(t)$.

Discrete Hermite functions. We call the vector

$$\Phi_\ell = (\varphi_\ell(\alpha_0, 1) \quad \varphi_\ell(\alpha_1, 1) \quad \dots \quad \varphi_\ell(\alpha_{N-1}, 1))^T \quad (23)$$

the ℓ th discrete Hermite function.

As previously, we select t_k in (22) so that the length of intervals $[t_{k-1}, t_k]$ is constant⁴: $t_k - t_{k-1} = \Delta$. In this case the computation of each coefficient c_ℓ in (22) can be seen as the inner product of a discrete signal

$$\mathbf{s} = (s(\alpha_0 \lambda) \quad s(\alpha_1 \lambda) \quad \dots \quad s(\alpha_{2K+1} \lambda))^T$$

with the ℓ th discrete Hermite function Φ_ℓ

$$c_\ell \approx \Delta \sum_{k=-K}^K s(\alpha_{k+K} \lambda) \varphi_\ell(\alpha_{k+K}, 1) = \Delta \cdot \langle \Phi_\ell, \mathbf{s} \rangle.$$

Expansion matrix. The matrix Φ in (13) can now be rewritten as

$$\Phi = \pi^{-\frac{1}{4}} W \mathcal{P}_{P,\alpha} \quad (24)$$

where

$$W = \text{diag}(e^{-\frac{\alpha_k^2}{2}})_{0 \leq k < 2K+1}$$

is a diagonal matrix, and $\mathcal{P}_{P,\alpha}$ is given in (17).

Finally, if $M = 2K + 1$, and Φ is a square $(2K+1) \times (2K+1)$ matrix, then it follows from (19) that the rows of Φ form an orthogonal basis:

$$\Phi \Phi^T = \pi^{-\frac{1}{2}} W^2 D^{-1}. \quad (25)$$

We denote the diagonal matrix on the right-hand side (RHS) of (25) as Λ . Later, to account for the vector norms, we will pre-multiply the input signal \mathbf{s} with the weight matrix Λ^{-1} .

⁴To our best knowledge, there is no closed-form expression for the roots of Hermite polynomials $P_{2K+1}(t)$ [10], [13]. In general, we cannot guarantee that for any K there exists λ , such that $\lambda \alpha_{k+K}$ can be placed inside intervals $[t_{k-1}, t_k]$ of equal length, so that $t_{k-1} \leq \lambda \alpha_{k+K} \leq t_k$ holds. However, in the experiments in Section V only $K \leq 15$ were used, and we numerically confirmed the existence of suitable values λ . Moreover, for each considered K there exists a range of values of λ that satisfy the condition; in the experiments, the value that maximized the compression ratio was selected.

Proposed algorithm. The proposed compression algorithm operates as follows. We assume that the ECG signal $s(t)$ is sampled at points $\alpha_{k+K}\lambda$, $-K \leq k \leq K$, to obtain a vector of samples

$$\mathbf{s} = (s(\alpha_0\lambda), s(\alpha_1\lambda), \dots, s(\alpha_{2K}\lambda))^T. \quad (26)$$

The selection of parameter λ is discussed in Section V. Then we compute the vector of expansion coefficients

$$\mathbf{c} = \Delta\Phi^T\Lambda^{-1}\mathbf{s} \quad (27)$$

where Φ and Λ are given in (24) and (25). Following this, we construct the vector $\hat{\mathbf{c}}$ by keeping only M coefficients in \mathbf{c} with the largest magnitudes and setting others to zero. Finally, we use $\hat{\mathbf{c}}$ to obtain the signal approximation

$$\hat{\mathbf{s}} = \Delta^{-1}\Phi\hat{\mathbf{c}}. \quad (28)$$

Advantages. The proposed algorithm addresses several limitations of the original compression algorithms based on continuous Hermite functions.

First of all, the discrete Hermite functions Φ_ℓ in (23) form an orthogonal basis in the finite-dimensional vector space \mathbb{C}^N of all discrete signals of length N . The expansion of QRS complexes sampled at N nodes into the discrete Hermite functions is complete, and allows for a perfect reconstruction of the sampled complexes. By contrast, an expansion of a continuous QRS complex into a finite number $N < \infty$ of continuous Hermite functions $\varphi_\ell(t, \sigma)$ is not complete, since these N functions do not form a complete basis in $\mathcal{L}^2(\mathbb{R})$. As a result, a perfect reconstruction cannot be achieved.

Furthermore, for a digital implementation of the compression, both the QRS complexes and the basis functions must be discretized by sampling. As we have discussed, equidistant sampling of continuous Hermite functions does not yield an orthogonal basis in the \mathbb{C}^N space. In this case, approximating a signal with more basis functions may not necessarily lead to a smaller approximation error. On the other hand, since the discrete Hermite functions form an orthogonal basis, increasing the number of vectors used for the signal approximation necessarily decreases the approximation error [24]. Moreover, since we pre-compute all coefficients c_0, c_1, \dots, c_{2K} , and only after this select M ones with the largest magnitude to obtain $\hat{\mathbf{c}}$, the approximation error is minimized for a fixed M . In addition, if an exact reconstruction of signals is required, it can be achieved by using $M = 2K + 1$ coefficients to obtain $\hat{\mathbf{c}}$.

Finally, the proposed algorithm has a more efficient implementation, since the computational cost of Φ in (24) (as well as Φ^T) is approximately two times lower compared to the cost of Φ (and Φ^T) in (13), as discussed in Section III-B.

V. EXPERIMENTS

Setup. In order to analyze the performance of the proposed compression algorithm, we study the compression of QRS complexes extracted from ECG signals in the MIT-BIH ECG Compression Test Database [25]. The database contains 168 ECG

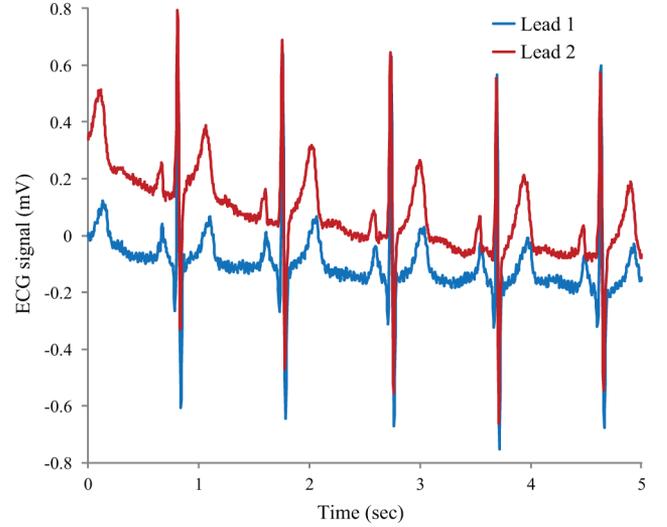


Fig. 3. An example ECG signal with two leads.

signals sampled at 250 Hz. An example ECG signal with two recorded channels, or *leads*, is shown in Fig. 3.

Preprocessing. To make the experiments uniform across the database, we only used the first 10 seconds of the first lead of each signal. All detectable QRS complexes were extracted automatically, a total of $N = 1486$ complexes. During the extraction, we require that each extracted complex is centered around the R peak. As a result, each extracted complex is available as a discrete signal of length $2K + 1 \in \{27, 29, 31\}$, where the $(K + 1)$ th sample corresponds to the R peak. Hence, each signal represents a continuous QRS complex of duration 104, 112, or 120 milliseconds sampled at 250 Hz.

For the new compression algorithm, QRS complexes must be sampled at points $\alpha_{k+K}\lambda$ proportional to the roots of $P_{2K+1}(t)$. These signals are not available directly from MIT-BIH ECG Compression Test Database. Hence, we had to construct the required discrete signals prior to running the experiments. We first reconstructed continuous QRS complexes by interpolating the extracted discrete QRS complexes with *sinc* functions, and then sampled the obtained continuous signals at the required points $\alpha_{k+K}\lambda$. The value of the parameter λ was determined experimentally for each lead to maximize the resulting compression ratio. In particular, we have observed that the optimal value of λ “stretches” the weighted signal $\Lambda^{-1}\mathbf{s}$ in (27) to make its shape closely resemble the shapes of the discrete Hermite functions Φ_ℓ in (23). Since in our experiments having only a few coefficients $c_\ell = \langle \Lambda^{-1}\mathbf{s}, \Phi_\ell \rangle$ with large magnitudes implies a higher compression ratio, we hypothesize that the optimal choice of λ may be explained by the Cauchy-Schwarz inequality [26] that states that the magnitude of the inner product of two vectors with equal norms is maximized when the vectors are equal. Thus, if $\Lambda^{-1}\mathbf{s}$ closely resembles Φ_ℓ , their inner product $c_\ell = \langle \Lambda^{-1}\mathbf{s}, \Phi_\ell \rangle$ should have a large magnitude.

Similarly, we determined the value of the parameter σ for each lead to maximize the compression ratio for the original compression algorithm that expands QRS complexes into continuous Hermite functions.

TABLE I
AVERAGE COMPRESSION RATIOS OF DIFFERENT ALGORITHMS FOR 10%, 15%, 20%, AND 25% APPROXIMATION ERRORS. THE AVERAGE NUMBER OF COEFFICIENTS REQUIRED TO ACHIEVE THE DESIRED ERROR IS INDICATED IN THE PARENTHESIS NEXT TO THE CORRESPONDING RATIO

Error	Proposed algorithm	Original algorithm	DFT-based	DCT-based	DWT-based
10%	5.3 (5.8)	3.5 (9.0)	3.7 (8.3)	4.3 (7.3)	3.3 (9.4)
15%	7.0 (4.4)	4.3 (7.2)	4.2 (7.4)	5.1 (6.1)	4.2 (7.5)
20%	9.2 (3.4)	5.0 (6.2)	4.6 (6.7)	5.8 (5.3)	4.8 (6.5)
25%	10.4 (2.9)	5.8 (5.4)	5.1 (6.1)	6.6 (4.7)	5.5 (5.6)

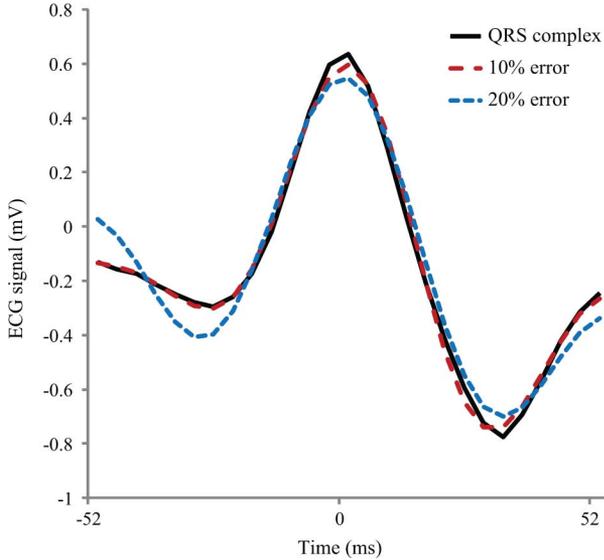


Fig. 4. A QRS complex and its approximations with 10% and 20% errors.

In a real-time system, the parameters λ or σ would be pre-set for each lead. In addition, the incoming continuous ECG signal would be immediately sampled at the desired points. Hence, the interpolation step would not be required.

Objective. The main objective of the experiments is to achieve the maximal compression ratio for a desired approximation error. We define the error as the norm of the difference between the original signal and its approximation, normalized by the norm of the original signal

$$\text{approximation error} = \frac{\|\hat{s} - s\|_2}{\|s\|_2}. \quad (29)$$

In this paper, we seek to achieve a 10% approximation error. It has been verified by one of the authors who is an experienced cardiologist, that all signals in the MIT-BIH ECG Compression Database can be compressed with the methods considered in this paper with the 10% error while remaining *medically acceptable*. This means that visually all signals remain sufficiently similar to their original versions, and would not lead to incorrect interpretation and diagnoses.

We also provide compression ratios for 15%, 20%, and 25% approximation errors. In general, signals in the MIT-BIH ECG Compression Database compressed with these errors have been found medically unacceptable. For example, in our experiments for the approximation error of 20%, 630 out of 1486 (42.4%) compressed QRS complexes were identified as medically unacceptable. However, most of the distortions were introduced at

the boundaries of the compressed QRS complexes, as demonstrated in the example shown in Fig. 4. Potentially, this problem could be addressed by smoothing the boundaries, and hence reducing the approximation error.

Compression algorithms. For the original compression algorithm that expands QRS complexes into continuous Hermite functions, as described in Section II-B, we compute a sufficient number M of coefficients c_0, \dots, c_{M-1} in (11). We use the minimal possible M that ensures that the reconstruction \hat{s} in (12) has the desired approximation error (29).

For the new compression algorithm, we compute $2K + 1$ coefficients, and select a sufficient number M of the coefficients with the largest magnitude to obtain the approximation \hat{s} in (28) that yields the required approximation error (29).

In addition, we study the accuracy of compression algorithms based on widely used orthogonal discrete signal transforms. In particular, we consider the discrete Fourier transform (DFT), the discrete cosine transform (DCT), and discrete wavelet transform (DWT). In the latter case, we consider an orthogonal DWT based on Daubechies filters of length 4 with three levels of decomposition [24]. Since all these transforms are orthogonal, we can replace Φ in (11) and (12) with the corresponding transform, apply the transforms to the signal s in (10), and select a sufficient number M of coefficients with the largest magnitudes, such that reconstruction \hat{s} in (12) has the desired approximation error (29).

Results. The average compression ratios for 10%–25% approximation errors are shown in Table I. The ratios were computed as

$$\text{average ratio} = \frac{\sum_{n=0}^{N-1} (2K_n + 1)}{\sum_{n=0}^{N-1} M_n}.$$

Here, $2K_n + 1 \in \{27, 29, 31\}$ is the length of the n th sampled QRS complex, and M_n is the number of coefficients required to achieve the desired approximation error. Naturally, the higher the compression ratio, the better the algorithm performance, since fewer coefficients are required.

In Table I, we also identify the average minimal number M of coefficients c_{ℓ_m} required to achieve the desired approximation error. This characteristic is included for implementation purposes and the ease of interpretation of compression ratios.

VI. DISCUSSION AND CONCLUSIONS

As we observe from Table I, the proposed compression algorithm has the highest compression ratio for all considered approximation errors. In particular, it requires on average only 6 coefficients to reconstruct compressed QRS complexes that are

medically acceptable. This is a 25% improvement compared to using DCT as the compressing transform, a 33% improvement compared to using DFT or the compression algorithm based on continuous Hermite functions, and a 40% improvement compared to using DWT. As an interesting observation, recall that here we are using the DWT with three decomposition levels; we tested DWT with other numbers of decomposition levels, and the compression ratios were even lower.

The larger compression ratio of the proposed algorithm in comparison with the original expansion into continuous Hermite functions can be explained by the improvements discussed in Sections II and IV. The better performance of the proposed algorithm compared to the orthogonal transforms DFT, DCT, and DWT, on the other hand, may have different reasons. One potential factor is that, as discussed above, the proposed compression algorithm uses the parameter λ to better “fit” the input signals to the basis of discrete Hermite functions that have shapes resembling QRS complexes. This “fitting” may lead to fewer coefficients with large magnitudes compared to DFT, DCT, and DWT, which have corresponding basis functions of different shapes. Another, more speculative, potential factor is that DFT, DCT, and DWT are well suited for the compression of “smooth” signals (for the discussion, see [24] and references therein), since they efficiently capture the low-frequency components of signals using a small fraction of coefficients, and remove the high-frequency component that may require many additional coefficients. Since some QRS complexes may have significant high-frequency components, they may require a larger number of coefficients to achieve a desired reconstruction accuracy. In our experiments, however, the effect of high frequencies in QRS complexes on the compression performance was observed only in a limited number of cases.

Future improvements. As we mentioned in Section V, the compression with an approximation error higher than 10% frequently leads to medically unacceptable signals. However, most of the distortion is introduced at the boundaries of the compressed QRS complexes.

Fig. 4 shows an example of such distortions for the 20% approximation error. While the shape of the compressed signal is similar to the original QRS complex, the compressed signal is clearly corrupted at the left and right boundaries. As a result, discontinuities are introduced on the edge between the compressed QRS complex and the preceding P-R segment, as well as on the edge of the QRS complex and the following S-T segment. Potentially, one could address this problem by smoothing the compressed signal at the boundaries, thus reducing the approximation error.

Conclusions. We have constructed a novel algorithm for the compression of QRS complexes. The proposed algorithm is based on the expansion of signals with compact support (such as ECG signals) into the basis of discrete Hermite functions. These functions are constructed by sampling the continuous Hermite functions at sampling points proportional to the roots of a corresponding Hermite polynomial.

The proposed algorithm uses results from our recently developed theory of signal models for orthogonal polynomials. As confirmed by the experiments, the novel algorithm achieves a higher compression ratio compared with the original algorithm

based on the expansion into continuous Hermite functions, as well as other widely used compression algorithms.

In addition, we developed a fast computational algorithm for the proposed compression method. The proposed approach reduces the number of operations approximately by a factor of two.

REFERENCES

- [1] P. Libby, R. O. Bonow, D. L. Mann, and D. P. Zipes, *Braunwald's Heart Disease: A Textbook of Cardiovascular Medicine, Single Volume*, 8th ed. New York: Saunders, 2007.
- [2] A. Sandryhaila, J. Kovacevic, and M. Püschel, “Compression of QRS complexes using Hermite expansion,” in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2011, pp. 581–584.
- [3] M. Püschel and J. M. F. Moura, “Algebraic signal processing theory” [Online]. Available: <http://arxiv.org/abs/cs.IT/0612077>
- [4] A. Sandryhaila, “Algebraic Signal Processing: Modeling and Subband Analysis,” Ph.D., Carnegie Mellon Univ., Pittsburgh, PA, 2010.
- [5] A. Sandryhaila, J. Kovacevic, and M. Püschel, “Algebraic signal processing theory: 1-D nearest-neighbor models,” *IEEE Trans. Signal Process.*, submitted for publication.
- [6] P. Laguna, R. Jané, S. Olmos, N. V. Thakor, H. Rix, and P. Caminal, “Adaptive estimation of QRS complex wave features of ECG signal by the Hermite model,” *J. Med. Biol. Eng. Comput.*, vol. 34, no. 1, pp. 58–68, 1996.
- [7] L. Sörnmo, P. O. Börjesson, P. Nygård, and O. Pahlm, “A method for evaluation of QRS shape features using a mathematical model for the ECG,” *IEEE Trans. Biomed. Eng.*, vol. BME-28, no. 10, pp. 713–717, 1981.
- [8] L. R. L. Conte, R. Merletti, and G. V. Sandri, “Hermite expansion of compact support waveforms: Applications to myoelectric signals,” *IEEE Trans. Biomed. Eng.*, vol. 41, no. 12, pp. 1147–1159, 1994.
- [9] M. Lagerholm, C. Peterson, G. Braccini, L. Edenbrandt, and L. Sörnmo, “Clustering ECG complexes using Hermite functions and self-organizing maps,” *IEEE Trans. Biomed. Eng.*, vol. 47, no. 7, pp. 838–848, 2000.
- [10] G. Szegő, *Orthogonal Polynomials*, 3rd ed. New York: Amer. Math. Soc. Colloq., 1967.
- [11] R. Askey, *Orthogonal Polynomials and Special Functions*. Philadelphia, PA: SIAM, 1966.
- [12] B. S. Kashin and A. A. Saakyan, *Orthogonal Series*. New York: Amer. Math. Soc., 1989.
- [13] W. Gautschi, *Orthogonal Polynomials: Computation and Approximation*. Oxford, U.K.: Oxford Univ. Press, 2004.
- [14] M. Püschel, “Cooley-Tukey FFT like algorithms for the DCT,” *Proc. ICASSP*, vol. 2, pp. 501–504, 2003.
- [15] M. Püschel and J. M. F. Moura, “The algebraic approach to the discrete cosine and sine transforms and their fast algorithms,” *SIAM J. Comput.*, vol. 32, no. 5, pp. 1280–1316, 2003.
- [16] M. Püschel and J. M. F. Moura, “Algebraic signal processing theory: Foundation and 1-D time,” *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3572–3585, 2008.
- [17] M. Püschel and J. M. F. Moura, “Algebraic signal processing theory: 1-D space,” *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3586–3599, 2008.
- [18] M. Püschel and J. M. F. Moura, “Algebraic signal processing theory: Cooley-Tukey type algorithms for DCTs and DSTs,” *IEEE Trans. Signal Process.*, vol. 56, no. 4, pp. 1502–1521, 2008.
- [19] J.-B. Martens, “The Hermite transform—Theory,” *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 9, pp. 1595–1605, 1990.
- [20] J.-B. Martens, “The Hermite transform—Applications,” *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 9, pp. 1607–1618, 1990.
- [21] A. Jerri, *Integral and Discrete Transforms with Applications and Error Analysis*. Boca Raton, FL: CRC, 1992.
- [22] J. R. Driscoll, D. M. Healy, Jr., and D. Rockmore, “Fast discrete polynomial transforms with applications to data analysis for distance transitive graphs,” *SIAM J. Comput.*, vol. 26, pp. 1066–1099, 1997.
- [23] A. Sandryhaila, J. Kovacevic, and M. Püschel, “Algebraic signal processing theory: Cooley-Tukey type algorithms for polynomial transforms based on induction,” *SIAM J. Matrix Analysis Appl.*, vol. 32, no. 2, pp. 364–384, 2011.
- [24] S. Mallat, *A Wavelet Tour of Signal Processing: The Sparse Way*, 3rd ed. New York: Academic, 2008.

- [25] MIT-BIH ECG Compression Test Database [Online]. Available: <http://www.physionet.org/physiobank/database/cdb>
- [26] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York: Dover, 1972.



Aliaksei Sandryhaila (S'06–M'10) received the B.S. degree in computer science from Drexel University, Philadelphia, PA, in 2005, and the Ph.D. degree in electrical and computer engineering from Carnegie Mellon University (CMU), Pittsburgh, PA, in 2010.

He is currently a Postdoctoral Researcher with the Department of Electrical and Computer Engineering, CMU, and a Senior Research Scientist at SpiralGen, Inc. His research interests include signal processing, applied mathematics, medical signal processing, data modeling and analysis, and fast algorithm design.



Samir Saba received the B.S. degree in electrical engineering and the M.D. degree from the American University of Beirut in 1985 and 1993, respectively.

He completed his training in internal medicine (1993–1996), in cardiovascular diseases (1996–1999), and in cardiac electrophysiology (1999–2000) at Tufts University—New England Medical Center, Boston, MA. He then joined the faculty of the University of Pittsburgh Medical Center, where he is now an Associate Professor of Medicine in the field of cardiac electrophysiology.

Since 2005, he has been the Director of the Cardiac Electrophysiology section at the University of Pittsburgh. He has authored more than 100 manuscripts that were published in peer-reviewed journals. He is also a busy clinical electrophysiologist with expertise in cardiac device implantations and ablation procedures. His research interests include cardiac device therapy for heart failure, signal processing of intracardiac electrical signals for ischemia detection, as well as the study of the molecular and genetic bases of arrhythmias.

Dr. Saba is currently a Fellow of the American College of Cardiology and the Heart Rhythm Society. He is the recipient of research grants from the National Institutes of Health, the American Heart Association, the American Heart Foundation, and the American College of Cardiology.



Markus Püschel (M'99–SM'05) received the Diploma (M.Sc.) in mathematics and the Doctorate (Ph.D.) degree in computer science, in 1995 and 1998, respectively, both from the University of Karlsruhe, Germany.

He is a Professor of Computer Science at ETH Zurich, Switzerland. Before, he was a Professor of Electrical and Computer Engineering at Carnegie Mellon University, Pittsburgh, PA, where he still has an adjunct status. From 1998 to 1999, he was a Postdoctoral Researcher at Mathematics and Computer

Science, Drexel University, Philadelphia, PA. From 2000 to 2010, he was with Carnegie Mellon University, and since 2010 he has been with ETH Zurich. In 2009, he cofounded SpiralGen, Inc.

Dr. Püschel was an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING, the IEEE SIGNAL PROCESSING LETTERS, was a Guest Editor of the PROCEEDINGS OF THE IEEE and the *Journal of Symbolic Computation*, and served on various program committees of conferences in computing, compilers, and programming languages. He is a recipient of the Outstanding Research Award of the College of Engineering at Carnegie Mellon and the Eta Kappa Nu Award for Outstanding Teaching. He also holds the title of Privatdozent at the University of Technology, Vienna, Austria.



Jelena Kovacevic (S'88–M'91–SM'96–F'02) received the Dipl. Electr. Eng. degree from the Electrical Engineering Department, University of Belgrade, Yugoslavia, in 1986, and the M.S. and Ph.D. degrees from Columbia University, New York, in 1988 and 1991, respectively.

She is a Professor of Biomedical Engineering and Electrical and Computer Engineering and the Director of the Center for Bioimage Informatics, Carnegie Mellon University, Pittsburgh, PA, which she joined in 2003. From 1991 to 2002, she was

with Bell Labs, Murray Hill, NJ. She was a cofounder and Technical VP of xWaveforms, based in New York City. She was also an Adjunct Professor at Columbia University. Her research interests include bioimaging, as well as multiresolution techniques, such as wavelets and frames. She is a coauthor (with M. Vetterli) of the book *Wavelets and Subband Coding* (Englewood Cliffs, NJ: Prentice-Hall, 1995).

Dr. Kovacevic coauthored a top-10 cited paper in the *Journal of Applied and Computational Harmonic Analysis*, and the paper for which A. Mojsilovic received the Young Author Best Paper Award. Her paper on multidimensional filter banks and wavelets (with M. Vetterli) was selected as one of the Fundamental Papers in Wavelet Theory. She received the Belgrade October Prize in 1986 and the E.I. Jury Award at Columbia University in 1991. She was the Editor-in-Chief of the IEEE TRANSACTIONS ON IMAGE PROCESSING. She served as an Associate Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING, as a Guest Co-Editor (with I. Daubechies) of the Special Issue on Wavelets of the PROCEEDINGS OF THE IEEE, Guest Co-Editor (with M. Vetterli) of the Special Issue on Transform Coding of the IEEE SIGNAL PROCESSING MAGAZINE, and Guest Co-Editor (with R. F. Murphy) of the Special Issue on Molecular and Cellular Bioimaging of the IEEE SIGNAL PROCESSING MAGAZINE. She has been on the Editorial Boards of the *Foundations and Trends in Signal Processing*, SIAM book series on Computational Science and Engineering, *Journal of Applied and Computational Harmonic Analysis*, *Journal of Fourier Analysis and Applications*, and the IEEE SIGNAL PROCESSING MAGAZINE. She is a regular member of the NIH Microscopic Imaging Study Section. From 2000 to 2002, she served as a Member-at-Large of the IEEE Signal Processing Society Board of Governors. She is the Chair of the Bio Imaging and Signal Processing Technical Committee. She was the General Chair of ISBI'06, General Co-Chair (with V. Goyal) of the DIMACS Workshop on Source Coding and Harmonic Analysis and General Co-Chair (with J. Allebach) of the Ninth IMDSP Workshop. She was a plenary/keynote speaker at the 20 Years of Wavelets'09, European Women in Mathematics'09, MIAABWorkshop'07, Statistical Signal Processing Workshop'07, Wavelet Workshop '06, NORISIG '06, ICIAR '05, Fields Workshop '05, DCC '98, as well as SPIE '98.