

Is the fast Hankel transform faster than quadrature?

Kerry Key¹

ABSTRACT

The fast Hankel transform (FHT) implemented with digital filters has been the algorithm of choice in EM geophysics for a few decades. However, other disciplines have predominantly relied on methods that break up the Hankel transform integral into a sum of partial integrals that are each evaluated with quadrature. The convergence of the partial sums is then accelerated through a nonlinear sequence transformation. While such a method was proposed for geophysics nearly three decades ago, it was demonstrated to be much slower than the FHT. This work revisits this problem by presenting a new algorithm named quadrature-with-extrapolation (QWE). The QWE method recasts the quadrature sum into a form conceptually similar

to the FHT approach by using a fixed-point quadrature rule. The sum of partial integrals is efficiently accelerated using the Shanks transformation computed with Wynn's ϵ algorithm. A Matlab implementation of the QWE algorithm is compared with the FHT method for accuracy and speed on a suite of relevant modeling problems including frequency-domain controlled-source EM, time-domain EM, and a large-loop magnetic source problem. Surprisingly, the QWE method is faster than the FHT for all three problems. However, when the integral needs to be evaluated at many offsets and the lagged convolution variant of the FHT is applicable, the FHT is significantly faster than the QWE method. For divergent integrals such as those encountered in the large loop problem, the QWE method can provide an accurate answer when the FHT method fails.

INTRODUCTION

Electromagnetic (EM) modeling often relies on evaluating integrals of the form

$$F(r) = \int_0^{\infty} f(k)g(kr)dk, \quad (1)$$

where $g(kr)$ is an oscillatory Bessel function or sinusoid. The term $f(k)$ is the kernel function that depends on the subsurface physical properties and may also be oscillatory. Due to the oscillatory behavior of $g(kr)$, standard quadrature methods applied to these integrals can be slow to converge or may fail if the integral is divergent. Consequently, special care is required for their numerical evaluation.

The ubiquity of these integrals in EM geophysics motivates the need for accurate and efficient numerical integration techniques. Consider the Hankel transform integral:

$$F(r) = \int_0^{\infty} f(\lambda)J_i(\lambda r)d\lambda, \quad (2)$$

where J_i is an i th order Bessel function of the first kind that arises from the cylindrical symmetry of the problem. This integral may need to be evaluated millions of times or more for 3D modeling applications, where 1D solutions are used to generate the source terms required by scattered-field formulations (e.g., Newman and Alumbaugh, 1995), or are used for the Green's tensors required by integral equation methods (e.g., Raiche, 1999). Numerous evaluations are also necessary for 1D inversions. For example, Bayesian inversion typically requires 10,000–100,000 or more forward computations (e.g., Chen et al., 2007). The sine and cosine transforms used by 2.5D EM algorithms (e.g., Li and Key, 2007) and time-domain EM (e.g., Newman et al., 1986) are other examples of heavily used oscillatory integrals.

Most EM geophysics modeling codes in use today evaluate these integrals by using the digital filter method that was proposed by Ghosh (1971). Briefly, the digital filter method can be found by substituting $r = e^x$ and $\lambda = e^{-y}$ into equation 2, giving

Peer-reviewed code related to this article can be found at <http://software.seg.org/2012/0003>.

Manuscript received by the Editor 1 July 2011; published online 5 April 2012.

¹University of California at San Diego, Scripps Institution of Oceanography, La Jolla, California, USA. E-mail: kkey@ucsd.edu.

© 2012 Society of Exploration Geophysicists. All rights reserved.

$$e^x F(e^x) = \int_{-\infty}^{\infty} f(e^{-y}) J_i(e^{x-y}) e^{x-y} dy, \quad (3)$$

which is then recast in the form of the convolution integral

$$I(x) = \int_{-\infty}^{\infty} f(y) h(x-y) dy = \int_{-\infty}^{\infty} f(x-y) h(y) dy. \quad (4)$$

The discrete approximation to the convolution is

$$I(j) = \sum_{i=-\infty}^{\infty} f(j-i) h(i), \quad (5)$$

where i and j are indices. Ghosh (1971) recognized that h is essentially a vector of linear filter coefficients that could be predetermined and subsequently applied to arbitrary kernel functions f . Optimal filter coefficients for finite length filters can be found by solving equation 5 for a length n filter using known integral transform pairs for I and f , where the transform pairs are selected to have rapid decay characteristics that are similar to the kernels of interest (e.g., Anderson, 1979). The resulting n -point digital filter approximation to equation 3 is

$$rF(r) \approx \sum_{i=1}^n f(b_i/r) h_i, \quad (6)$$

where the logarithmically spaced filter abscissae are

$$b_i = \lambda_i r = e^{ai}, \quad i = -l, -l+1, \dots, l, \quad (7)$$

$l = (n-1)/2$, and a is the spacing coefficient. The kernel function f tends to be numerically expensive but only needs to be evaluated at the n filter abscissae. Because common filter lengths are tens to hundreds of points, the digital filter method is usually much faster than simple direct quadrature and therefore has been referred to as the fast Hankel transform (FHT) (Anderson, 1979; Johansen and Sorensen, 1979). The simplicity and speed of the FHT, as well as the availability of published algorithms (e.g., Anderson, 1979, 1989) have led to its widespread use in EM geophysics.

Despite these advantages, the FHT has received little attention outside of geophysics. Instead, the engineering and applied math disciplines have largely relied on direct methods that break up the integral into an infinite sum of partial integrals which are each evaluated using quadrature. The sum of partial integrals is then accelerated or extrapolated to convergence through a suitable non-linear sequence transformation. See Lucas and Stone (1995) and Michalski (1998) for reviews of the various techniques available, which collectively can be referred to as integration, summation, and extrapolation (ISE) methods. Such numerical techniques date back at least to Longman (1956), who broke up the integral between the zeros of the Bessel function so that the sum of partial integrals forms a slowly converging alternating series, which was then accelerated through the use of Euler's transformation. Chave (1983) implemented a similar method for geophysical applications, where each partial integral is evaluated with adaptive quadrature (Patterson, 1968) and the sum of partial integrals is accelerated using a continued fraction expansion based on the method of Hänggi et al. (1978).

A significant benefit of ISE methods is that they can be computed iteratively until a specific error tolerance is obtained, where the error is estimated by differencing the results from successive iterations of the transformation. Conversely, the FHT does not have any intrinsic mechanism for estimating the error in the approximation to the integral. Additionally, ISE approaches can often handle divergent integrals that would otherwise cause the direct quadrature summation and FHT methods to fail.

Despite the error control advantages of the ISE approach, a subsequent GEOPHYSICS Discussion carried out between Chave and Anderson indicated that the FHT can obtain highly accurate results and, furthermore, is significantly faster than Chave's adaptive quadrature approach (Anderson, 1984; Chave, 1984). Consequently, the FHT method has dominated EM geophysics software, with little consideration given to alternative approaches in the ensuing decades. Anderson eventually released a hybrid algorithm that included both approaches, with the intention that the quadrature approach could be applied as a fail-safe when the FHT is insufficient, and also for independent testing of the digital filter accuracy (Anderson, 1989). Subsequent refinement of the FHT method has been carried out to handle problem-specific EM integrals (e.g., Sorensen and Christensen, 1994; Raiche, 1999) and new methods have been found for obtaining optimal digital filter coefficients (e.g., Guptasarma and Singh, 1997; Kong, 2007).

The disparity between the FHT being ubiquitous in EM geophysics while remaining largely absent in outside disciplines is intriguing. If the FHT is faster and acceptably accurate, why isn't it more prevalent in other disciplines that have similar Hankel transform integrals? A corollary to this question is why haven't the ISE methods heavily used in other disciplines been more thoroughly investigated for geophysical applications? These questions have motivated the present work, which presents a new ISE algorithm called quadrature-with-extrapolation (QWE). A simple Matlab implementation of the QWE algorithm is provided as a companion to this manuscript. This algorithm is compared with the FHT for accuracy and speed on a suite of relevant modeling problems including frequency-domain controlled-source EM, time-domain EM and a large-loop magnetic source problem. As shortly revealed, the results suggest that ISE methods deserve further consideration in EM geophysics.

QUADRATURE INTEGRATION WITH SEQUENCE EXTRAPOLATION

ISE methods partition the integral in equation 1 into an infinite sum of partial integrals according to:

$$F(r) = \int_0^{\infty} f(k)g(kr)dk = \sum_{i=0}^{\infty} F_i, \quad (8)$$

where

$$F_i = \int_{k_{i-1}}^{k_i} f(k)g(kr)dk. \quad (9)$$

The interval breakpoints $k_{i-1} < k_i$ are usually the zeros or extrema of the oscillatory function g , or some numerically convenient quantity with similar spacing. Because the direct sum in equation 8 usually converges slowly, extrapolation techniques are used to accelerate convergence. Weniger (1989) provides a thorough review

of many nonlinear sequence transformation techniques that have been proposed. Lucas and Stone (1995) and Michalski (1998) document the rapid convergence of various ISE methods applied to the Hankel transform integral of equation 2.

In the development of the present work, three transform methods were considered. Levin's t and u transformations were found to be effective, yet the older and simpler Shanks transformation implemented via the ϵ algorithm (Shanks, 1955; Wynn, 1956) was the only method of these three that could handle the divergent integral of the large-loop problem that is considered here. Furthermore, because the ϵ algorithm can be performed very efficiently, it was found to offer shorter run-times than the other transforms. The remainder of this work will therefore only consider the ϵ algorithm.

The ϵ algorithm is a recursive method defined by the three relations:

$$\epsilon_0^{(n)} = S_n, \quad n \geq 0, \quad (10)$$

$$\epsilon_1^{(n-1)} = \frac{1}{\epsilon_0^{(n)} - \epsilon_0^{(n-1)}}, \quad n \geq 1, \quad (11)$$

$$\epsilon_j^{(n-j)} = \epsilon_{j-2}^{(n-j+1)} + \frac{1}{\epsilon_{j-1}^{(n-j+1)} - \epsilon_{j-1}^{(n-j)}}, \quad n \geq 2, \quad 2 \leq j \leq n, \quad (12)$$

where S_n is the direct partial sum

$$S_n = \sum_{i=0}^n F_i. \quad (13)$$

This algorithm is graphically represented in Figure 1. The elements in the first column correspond to the direct sum of partial integrals, whereas elements to the right are the results of the sequence transformations. The n th order Shanks transformation is $\epsilon_n^{(0)}$ for even n and $\epsilon_{n-1}^{(1)}$ for odd n , as shown in Figure 1. If the partial sums S_n are regarded as part of a power series, the n th order Shanks transformation corresponds to the diagonal Padé approximants (Weniger, 1989); hence, the sequence transformation using the ϵ algorithm is equivalent to the continued fraction approach of Hänggi et al. (1978), yet as Michalski (1998) notes, the ϵ algorithm is more economical. We can therefore expect that ϵ algorithm will have the same rapid convergence demonstrated by the continued fraction approach of Chave (1983), yet will be more sparing in its computational run-time.

Lucas and Stone (1995) and Michalski (1998) examined the choice of breakpoints k_i , generally finding that the Bessel function zeros or extrema, or their approximations provide good convergence, with the best choice dependent on the specific problem.

The partial integrals in equation 9 are evaluated using a quadrature rule of the form

$$F_i \approx \sum_{j=1}^m w_j f(x_j/r) g(x_j), \quad (14)$$

where m is the quadrature order and w are weights associated with the quadrature abscissae x . Chave (1983) used the adaptive

Gauss-Kronrod quadrature method of Patterson (1968), in which the quadrature order is increased ($m = 3, 7, 15, 31, \dots$) until a predefined error criteria is satisfied. While this allows for a general code that can automatically adapt to the character of arbitrary kernel functions, there is a trade-off in the additional time required for the adaptive iterations. As is demonstrated here, a low-order quadrature rule is usually sufficient for many of the integrals encountered in EM geophysics. Consequently, this work will use a fixed-order Gauss quadrature method where the abscissae x are the roots of the corresponding Legendre polynomial (e.g., Trefethen, 2000).

A simple rearrangement of equation 14 reveals that the quadrature sum can be recast into a form remarkably similar to the FHT method in equation 6:

$$F_i \approx \sum_{j=1}^m f(x_j/r) w_j g(x_j) = \sum_{j=1}^m f(x_j/r) \hat{g}(x_j), \quad (15)$$

where $\hat{g}(x_j)$ is independent of the specific argument r . Hence, the product $w_j g(x_j)$ can be precomputed and stored for later use with any arbitrary kernel f , much like the precomputed filter weights h used in the digital filter method. This analogy can also be extended to the partial sum in equation 13. Defining the vector X as containing all the quadrature abscissae required for the n -th order sum:

$$X_n = [x_1, x_2, \dots, x_{mn}]^T, \quad (16)$$

the quadrature approximation to the sum in equation 13 can be written as the product of two vectors:

$$S_n = f(X/r)^T \hat{g}(X_n). \quad (17)$$

The term $\hat{g}(X_n)$ is relatively expensive to compute because it requires mn evaluations of the Bessel functions, but again this could be precomputed to exceedingly large n and stored in a subroutine, or it could be generated on the fly at the start of a program and then reused for arbitrary kernels and arguments r . This is analogous to the FHT in that the kernel function is evaluated at certain points and then multiplied by a weighting vector. The key difference here is

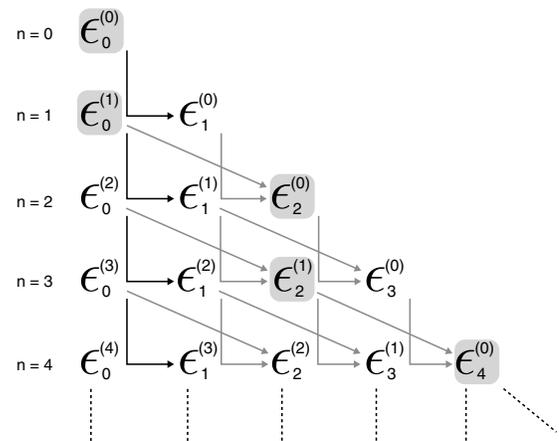


Figure 1. Illustration of the recursive ϵ algorithm. Black arrows represent the simple difference of equation 11, while gray arrows show the four-term recurrence relation of equation 12. Gray boxes show the n th order Shanks transformation (i.e., the best approximations) for odd and even orders n .

that the kernel function is evaluated at the unevenly spaced Gauss quadrature points rather than the logarithmically spaced abscissae of the digital filter.

This method of fixed-point quadrature with extrapolation via the Shanks transformation will be referred to as QWE. In the remainder of this work, a Matlab implementation of the QWE method is used to examine its speed and reliability for a suite of EM geophysics problems. The QWE method begins by computing the weighting vector $\hat{g}(X_n)$ for a large n (typically $n = 100$ is more than sufficient). This entails generating the quadrature weights and abscissae using the method described in Trefethen (2000) and the rapid computation of the Bessel function zeros using a vectorized implementation of Newton's method (e.g., Hamming, 1986). The integral of the first breakpoint interval is then evaluated with quadrature, yielding S_0 . The quadrature and extrapolation then proceeds iteratively as follows. For each iteration i , the quadrature sum F_i is computed with equation 15 and the direct sum is updated using $S_i = S_{i-1} + F_i$. The extrapolated sum is then found by performing the ϵ algorithm for order $n = i$ using equations 10–12. The Matlab implementation of these equations is patterned after the moving lozenge technique and Fortran subroutine given in Weniger (1989). Defining S_n^* as the extrapolated result from the ϵ algorithm, the absolute error δS_n^* is estimated by differencing subsequent iterations of the transform using

$$\delta S_n^* = |S_n^* - S_{n-1}^*|. \quad (18)$$

The QWE iterations proceed until the solution satisfies an user specified relative tolerance α and absolute tolerance β according to the inequality

$$\delta S_n^* \leq \alpha |S_n^*| + \beta. \quad (19)$$

NUMERICAL TESTS

For comparison purposes, the FHT was also implemented in Matlab. Digital filters of length 51, 101, and 201 points for the J_0 and J_1 Bessel functions and sine and cosine filters of length 101 and 201 points were designed using the approaches described in Guptasarma and Singh (1997) and Kong (2007). The FHT implementation follows the simple multiplication and sum given by equation 6. Anderson (1979) discusses a method known as adaptive convolution, where the sum is modified to be initially evaluated only over the center portion of the filter, and then followed out onto both tails until the resulting terms fall below some tolerance level, potentially reducing the number of costly kernel evaluations. Experience with the short filter lengths considered here has shown there is no appreciable time-savings from applying this adaptive method because most of the filter is usually required to minimize errors in the solution. Conversely, the 801 point filter of Anderson (1989) has large enough tails that some savings could be gained, depending on the specific integral under consideration.

The QWE method will now be compared to the FHT method for three different EM modeling problems. All of the Matlab codes and filter coefficients used for the following tests are included as a companion to this manuscript. These tests were performed with Matlab R2010b running on a Macbook Pro laptop with a 2.8 GHz Intel Core 2 Duo processor.

Controlled-source EM

During the past decade the controlled-source electromagnetic (CSEM) method has been adopted for offshore hydrocarbon exploration (e.g., Constable, 2010), making the marine CSEM problem a relevant test for the Hankel transform algorithms. For an inline horizontal electric dipole transmitter positioned in layer i of an arbitrary stack of N -layers with top depths z_i and conductivity σ_i , the frequency-domain horizontal electric field in the quasi-static approximation can be found with the transform:

$$E_h(r) = \int_0^\infty \left(\hat{E}_{J_0} J_0(\lambda r) + \hat{E}_{J_1} \frac{J_1(\lambda r)}{r} \right) d\lambda, \quad (20)$$

where the kernels are

$$\hat{E}_{J_0}(\lambda) = -\frac{1}{2\pi\mu\sigma_i} \left(\gamma_i^2 \hat{A}_h + \lambda^2 \frac{\partial \hat{\Lambda}_z}{\partial z} \right) \lambda, \quad (21)$$

$$\hat{E}_{J_1}(\lambda) = \frac{1}{2\pi\mu\sigma_i} \left(\hat{A}_h + \frac{\partial \hat{\Lambda}_z}{\partial z} \right) \lambda^2, \quad (22)$$

$\gamma_i^2 = \lambda^2 - i\omega\mu\sigma_i$, ω is the angular frequency and r is the horizontal offset between the transmitter source and the receiver location. When the source and receiver reside in the same layer i , the horizontal and vertical potentials have the forms (Key, 2009):

$$\hat{A}_h = a_i e^{\gamma_i(z-z_{i+1})} + b_i e^{-\gamma_i(z-z_i)} + \frac{\mu}{2\gamma_i} e^{-\gamma_i|z-z_s|}, \quad (23)$$

$$\begin{aligned} \hat{\Lambda}_z &= c_i e^{\gamma_i(z-z_{i+1})} + d_i e^{-\gamma_i(z-z_i)} \\ &\quad - \frac{\gamma_i}{\lambda^2} (a_i e^{\gamma_i(z-z_{i+1})} - b_i e^{-\gamma_i(z-z_i)}), \end{aligned} \quad (24)$$

where z is the measurement depth and z_s is the source depth. The coefficients a_i and b_i represent TE mode effects arising from conductivity contrasts below and above the source layer, respectively. Similarly, coefficients c_i and d_i represent TM mode effects. The multiplication by complex exponential terms accounts for decay from the layer boundaries to the measurement depth z . These coefficients vanish in a whole-space, leaving only the direct source term on the right side of equation 23. Recursive formulas for computing a_i , b_i , c_i , and d_i in an arbitrary stack of layers are given in the appendix of Key (2009).

In the parlance of Anderson (1982), \hat{E}_{J_0} and \hat{E}_{J_1} are called *related* kernels because they need to be evaluated concurrently and are computed using the same potentials given by equations 23 and 24. Similarly, the kernels for the other field components are also the progeny of these potentials. Hence, it is convenient to compute all of these transforms in parallel, which is done here for the QWE method and the digital filter method.

The kernels \hat{E}_{J_0} and \hat{E}_{J_1} could each be transformed independently by the ϵ algorithm. However, testing revealed that it is significantly more efficient to apply the transform to the sum of these terms. This is advantageous because the ϵ algorithm only needs to be performed on a single quantity rather than two, resulting in a time savings. Furthermore, it is sometimes the case that the integral of

one kernel is much larger than the other. By summing the terms, the lower magnitude term is numerically discounted during the sequence extrapolation such that the convergence is dictated only by the larger quantity. The partial integral breakpoints were set using the zeros of the J_1 Bessel function, which were found to be more efficient than the J_0 zeros, although the difference is insignificant compared to the differences observed between the QWE and FHT methods.

Figure 2 shows the \hat{E}_{J_0} and \hat{E}_{J_1} kernels computed at 1 Hz for a typical marine CSEM model known as the canonical reservoir model. This model consists of a resistive (10^{-12} S/m) air layer overlying a 1 km thick ocean of 3.3 S/m and a 1 S/m seafloor. A 100 m thick resistive hydrocarbon reservoir of 0.01 S/m is placed 1 km beneath the seafloor. The transmitter is positioned 10 m above the seabed. The kernels are fairly well behaved at small λ , but rapidly decay to zero when $\lambda \gg 1$.

The first test of the QWE method examines the convergence behavior using a relative tolerance of 10^{-12} and absolute tolerance of 10^{-30} . An example of the very rapid convergence is shown in Figure 3, where both kernel terms show a similar behavior, reaching the requested relative tolerance after 20 QWE iterations. For this example, the quadrature order was set to 51 to ensure a highly accurate solution. The next example examines the accuracy of significantly lower quadrature orders that offer increased numerical efficiency.

Figure 4 shows the accuracy of the responses computed for 21 seafloor receivers spaced from 0.5 to 20 km offset and for various quadrature orders. The baseline “truth” was estimated using the QWE method with 51-point quadrature and a relative tolerance of 10^{-12} . The relative error of the QWE approach rapidly decreases with increasing quadrature order, with order five taken as sufficient for most practical applications where 1% relative errors in real data are typically considered good. For $n = 9$ the relative error remains less than 10^{-5} , in excess of probably all practical needs. The relative errors are greatest at short ranges (<5 km). If further accuracy is desired, the order could be increased to 21 to drive the relative error down to negligible levels.

By performing a few simple tests such as this, the QWE method can be tailored to be accurate and fast for specific EM methods and

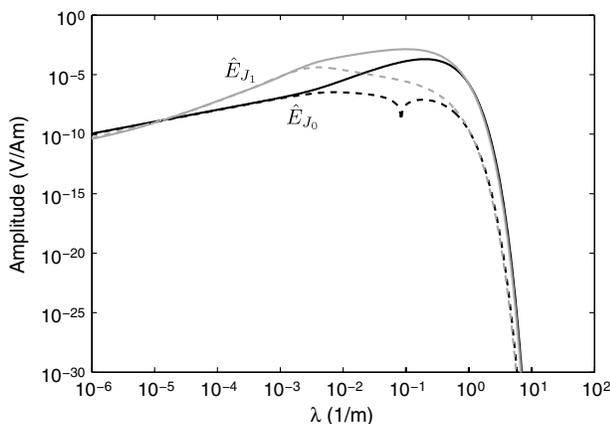


Figure 2. Horizontal electric field kernels from equations 21 and 22, computed for the canonical 1D model at 1 Hz. The real (solid) and the imaginary (dashed) components are shown as a function of the transform argument λ .

acquisition parameters. This is a significant benefit of ISE methods, because they can be adapted to meet the requested error level. Conversely, the digital filter method produces only a single result and the end-user is not provided with any estimation of its accuracy. Consider the 51-point digital filter response shown in the top of Figure 4. The smooth character of this response at long ranges offers

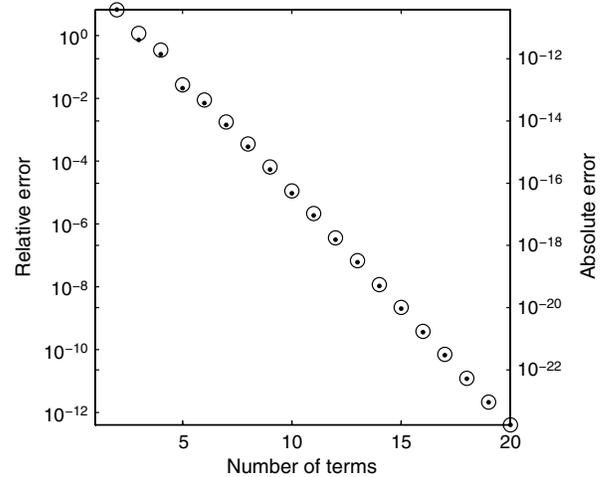


Figure 3. Convergence of the QWE method as a function of the order n shown for the relative error (circles) and absolute error (dots) for the inline electric field at 2 km offset.

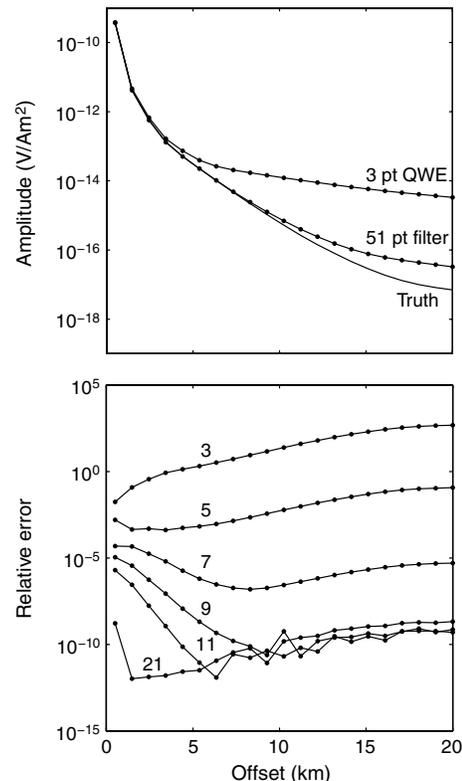


Figure 4. Top: Inline electric field responses computed for the canonical model using the QWE method with three-point quadrature and the FHT method with 51-point filters compared with the true CSEM response. Bottom: Relative error for the QWE method as a function of the numerical quadrature order.

no clue to the underlying fact that it is severely inaccurate. It also is fair to point out that a poor use of the QWE method will fail as well, as demonstrated with the very low-order three point quadrature. In both of these cases, an end-user is very likely to remain unaware that the solution is highly inaccurate; consequently, this anomalous large response might be interpreted as indicating a deep resistive structure that in fact does not exist.

Figure 5 shows comparisons between the QWE method and the FHT method using 51, 101, and 201 point filters. The relative error of the FHT using the very short 51-point filters exceeds 0.01 (1%) for most ranges, making it fairly worthless for most theoretical and practical applications. The 101-point filter performs better, yet exceeds 0.01 relative error at offsets great than 15 km. The 201-point filter performs the best, with errors below 10^{-6} for all ranges. The relative tolerance for the QWE method was set to 10^{-6} so that the results would be close to the accuracy of the 201-point filter. In general, the FHT results exhibit the smallest errors at short ranges, which then gradually increase at larger ranges. This is contrasted by

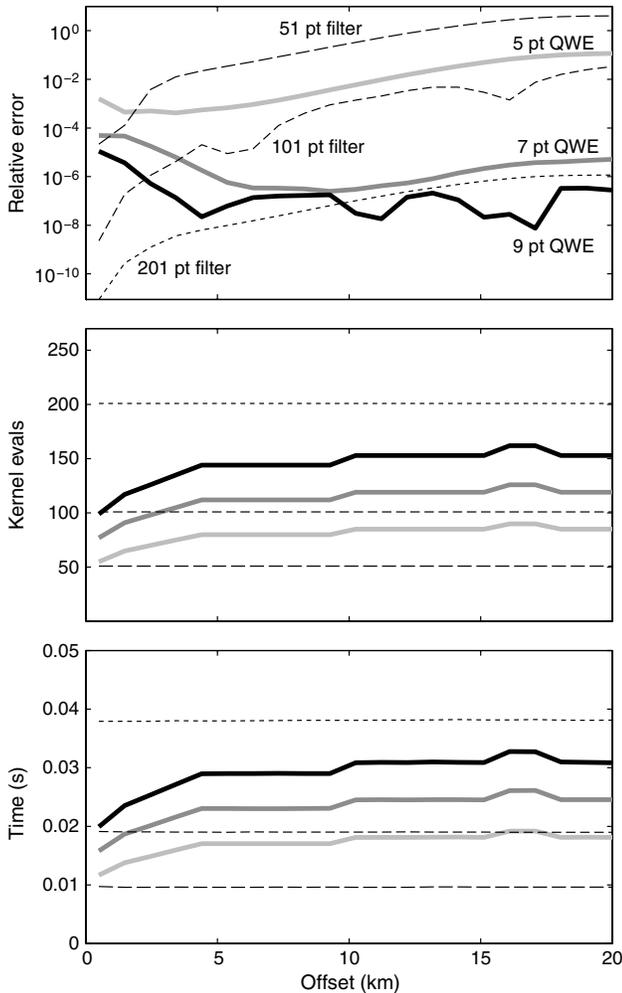


Figure 5. Comparison of the fast Hankel transform and the QWE method for marine CSEM computations. The top panel shows the relative error as a function of offset for 51, 101, and 201-point digital filters and for the QWE approach using five, seven, and nine-point quadrature. The middle panel shows the number of kernel evaluations required by each and the bottom panel shows the corresponding run-times required for each offset.

the QWE method, where the relative errors are largest at short ranges where the low quadrature orders are insufficient to meet the requested tolerance (yet in this example are sufficient for most practical needs).

Figure 5 shows that the QWE and FHT methods are competitive in terms of the number of kernel evaluations required. However, the QWE method clearly requires fewer kernel evaluations to reach a given accuracy. This is particularly evident for the nine-point QWE compared to the 201-point digital filter. The bottom of Figure 5 shows the CPU time required as a function of offset. Again, the QWE method performs favorably, particularly when compared to the 201-point filter. Because both methods use the same code to evaluate the kernel functions, the differences in run time reflect differences in the number of kernel evaluations as well as the time consumed by the typical 12–20 iterations of the ϵ algorithm required for the QWE method. Because the number of kernel evaluations are very similar for the 101-point filter and the seven-point QWE, the similar CPU times indicate that the ϵ algorithm is quite efficient.

Time-domain EM

This next example considers the time-domain impulse response, which can be obtained by Fourier transformation of the frequency domain results given in the previous section. Haines and Jones (1988) showed that the naive application of the fast Fourier transform to EM geophysics problems can require several tens of thousands of kernel evaluations to obtain an accurate response, and instead developed an efficient logarithmically sampled Fourier transform that only requires a few hundred kernel evaluations. However, it is common in the EM geophysics literature to invoke the causality of the impulse response so that it can be solved using either a cosine or sine transform

$$E_h(r, t) = \frac{2}{\pi} \int_0^{\infty} \text{Re}(E_h(r, \omega)) \cos(\omega t) d\omega, \quad (25)$$

$$= -\frac{2}{\pi} \int_0^{\infty} \text{Im}(E_h(r, \omega)) \sin(\omega t) d\omega, \quad (26)$$

which can be computed efficiently using the digital filter method (e.g., Newman et al., 1986).

Figure 6 shows an example of the kernel $E_h(r, \omega)$ at $r = 3$ km for the canonical reservoir model. At low frequency, the kernel is well-behaved, but at high frequencies, the kernel attenuates rapidly. The real and imaginary components contain significant oscillations, in contrast to relatively simple behavior of the wavenumber domain kernels shown in the previous section. The imaginary component has its first zero crossing at a higher frequency than the real component, making it favorable for numerical evaluation, hence the following study will focus on integrating the sine transform in equation 26. The QWE breakpoints are set to be the zeros of the sine function.

Time-domain responses are often required at time offsets of 0.01 to 100 s. The QWE method needs a slight modification to handle the earliest times. The first nonzero breakpoint corresponds to $\omega = \pi/t$, yet when $t \ll \pi$ the first quadrature interval will contain many kernel oscillations and hence a high quadrature order is required. Conversely, when $t \gg \pi$ the first interval and subsequent intervals will

only contain a small section of the kernel function and therefore can be accommodated with a low quadrature order. As simple solution to this complexity, the Matlab code for this problem inspects the kernel function to see if $E_h(r, \omega)$ decays by more than a factor of 100 in a given interval. If so, then Matlab's built-in adaptive Gauss-Kronrod quadrature routine `quadgk()` is called instead of the fixed-point rule. For the digital filter and QWE implementations of the sine transform, the kernel $E_h(r, \omega)$ is first obtained over a wide frequency range using the methods of the previous section, and then is interpolated using a piecewise cubic spline during the time-domain transform. A sampling density of 20 frequencies per decade was found to be sufficient for the model study given below.

Figure 7 shows a comparison of the QWE and digital filter methods for computing time-domain responses of the canonical model at 41 time steps spaced logarithmically in the range 0.01–100 s. Both methods perform well at late time offsets, but the 101-point filter responses break down the worst at early times while the QWE method and the 201-point filters produced accurate responses to lower amplitudes before hitting the early time numerical noise floor. However, the QWE approach obtains this accuracy in about half the run-time required by the 201-point digital filter.

Large-loop source

The two previous examples had kernel functions that decay rapidly at large arguments. This section considers a more difficult kernel that arises when using a large wire loop to generate a magnetic field. The transform for obtaining the vertical magnetic field in the plane of a horizontal loop at height h above a boundary is (e.g., Ward and Hohmann, 1987):

$$H_z(r) = \frac{r_t}{2} \int_0^{+\infty} \left[(R_{TE} e^{-2\gamma_1 h} + 1) \frac{\lambda^2}{\gamma_1} J_1(\lambda r_t) \right] J_0(\lambda r) d\lambda, \quad (27)$$

where r_t is the radius of the transmitter loop and r is the measurement range from the center of the transmitter loop. The term R_{TE} is the TE mode reflection coefficient, which for a two layer medium is

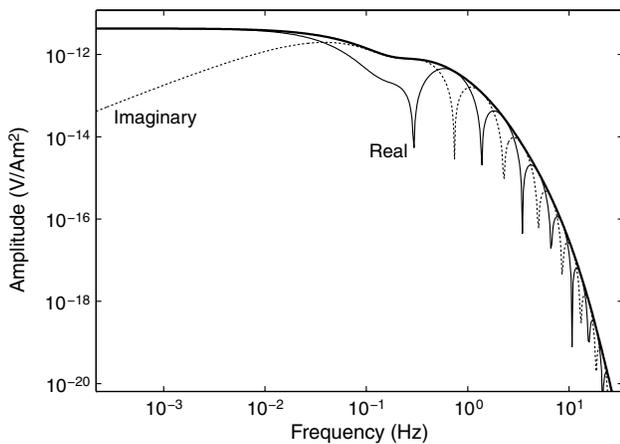


Figure 6. The time-domain kernel function $E(\omega, r)$ at $r = 3$ km range. The thick line shows the complex amplitude while the thin solid and dashed lines show the amplitude of the real and imaginary components.

$$R_{TE} = \frac{\gamma_1 \mu_2 - \gamma_2 \mu_1}{\gamma_1 \mu_2 + \gamma_2 \mu_1}, \quad (28)$$

where μ_1 and μ_2 are the relative magnetic permeabilities of each layer and the γ coefficient was defined earlier.

This transform is more complex because it contains the product of two Bessel functions and has a divergent term. Whereas the $R_{TE} e^{-2\gamma_1 h}$ term ultimately decays at large λ , the constant one multiplied by λ^2/γ_1 continues to increase with λ . Figure 8 shows an example of the real and imaginary parts of this kernel (without the Bessel functions), and the highly oscillatory Bessel functions. The divergent real kernel is proportional to λ^2 at small arguments and proportional to λ at large arguments (due to the relative importance of the frequency and conductivity terms in γ_1). Conversely, the imaginary kernel peaks at around $\lambda = 0.01 \text{ m}^{-1}$ and decays slowly in both directions.

For the QWE method, testing showed that using breakpoints based solely on the zeros of $J_1(\lambda r_t)$ is far more robust than using either $J_0(\lambda r)$ or the zeros of their product. For the FHT method, the kernel function is now ambiguous. For example, the quantity in brackets could be considered the kernel and the J_0 filter used, or the J_0 and J_1 terms could be interchanged and the J_1 filter used. Empirical testing showed the best (though, as soon revealed, insufficient) results were from using the J_1 filter when inside the loop at $r < r_t$ and the J_0 filter when $r \geq r_t$. For this problem, the QWE method used a relative error tolerance of 10^{-6} .

The results of the QWE and FHT methods are shown in Figure 9. Both methods performed well for the imaginary component, which can be explained by the imaginary kernel's well-behaved decay at large arguments. However, the FHT fails for the real component, giving a very inaccurate and obviously incorrect response. Conversely, the QWE method adequately handles this divergent kernel, capturing the singular response behavior near the transmitter radius

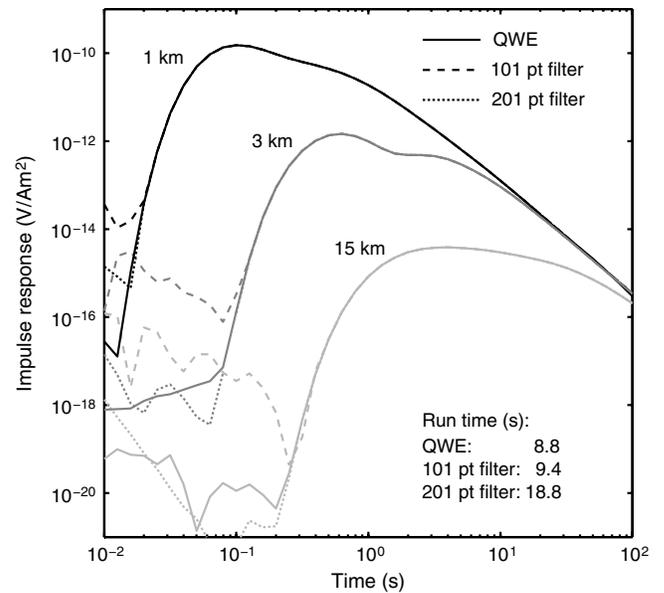


Figure 7. Time-domain impulse responses computed for the canonical 1D reservoir model at offsets r of 1, 3, and 15 km using the QWE and digital filter methods. The QWE method used a relative tolerance of 10^{-4} .

at $r = 5$ m. Furthermore, for most of the ranges considered, the QWE method required fewer kernel evaluations than the FHT.

While this problem serves as an example of the QWE method outperforming the FHT, it should be noted that when the measurement depth is sufficiently far from the plane of the source loop, the integral can become more tractable for the FHT because the constant one term in equation 27 is replaced by a decaying exponential term of the form $e^{-\gamma_1 d}$, where d is the vertical distance between the source loop and the measurement location.

Lagged-convolutions and spline approximations

This section examines how the QWE method compares to the lagged convolution variant of the FHT (Anderson, 1982). The lagged convolution enables the rapid solution at many offsets from a single source by leveraging the exponential spacing of the kernel evaluations, as briefly reviewed here. The digital filter used for the FHT requires the kernel to be evaluated at n discrete λ according to

$$\lambda_k = e^{ak}/r_0, \quad k = -l, -l+1, \dots, l, \quad (29)$$

where $l = (n-1)/2$, a is the filter spacing coefficient and r_0 is a particular offset. The nearby offset $r_1 = r_0 e^a$ can be computed with only a single additional kernel evaluation, as shown by:

$$\lambda_k = e^{ak}/r_1 = e^{a(k+1)}/r_0, \quad k = -l, -l+1, \dots, l. \quad (30)$$

This can be extended to a sequence of N offsets at the logarithmic spacing of the filter abscissae, all of which can be rapidly computed at the high accuracy of the digital filter using only $n + N - 1$ kernel evaluations, rather than the nN operations required if computed independently. The exponential spacing of the filter abscissae results in only a few additional kernel evaluations typically being required to cover a large range of offsets; hence, the lagged convolution scales very favorably with the number of offsets. Because in

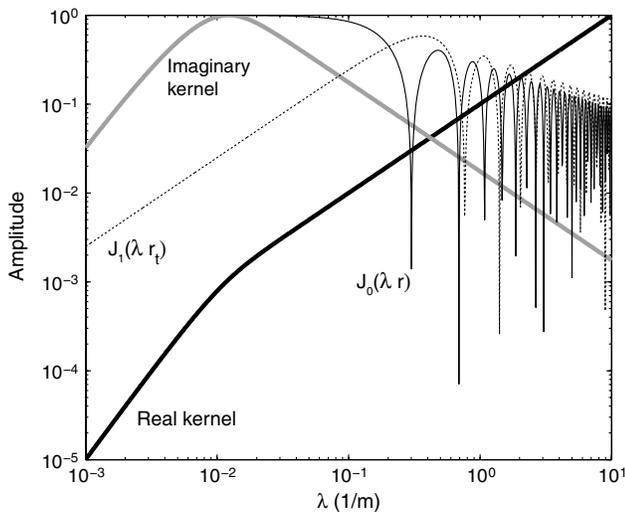


Figure 8. The kernel function for the large-loop problem (equation 27) evaluated for the parameters: $f = 25$ Hz, $r_t = 5$ m, $r = 8$ m, $h = 0$, $\sigma_1 = 10^{-12}$ S/m, $\sigma_2 = 1$ S/m, and $\mu_1 = \mu_2 = 1$. The real and imaginary parts have been normalized by their maximum values.

practice one usually needs the responses at arbitrary offsets, the lagged convolutions are first computed at the optimally spaced digital filter points that bound the arbitrary offsets and then these are interpolated to the specific offsets using cubic splines. The lagged convolution has great potential for very efficient computations, yet it should be noted that it can only be performed when all offsets have the same measurement depth and the same source position and orientation (i.e., they are computed from the same kernel function parameters).

The QWE method can also be sped up through spline interpolation. The kernel can be computed over a range of λ and then interpolated to obtain the particular λ required for the quadrature operations. However, because the quadrature and extrapolation operations are still required for each particular offset, the scaling with the number of offsets will be poorer than for the lagged convolution.

Figure 10 shows a comparison of the QWE and FHT methods applied to the canonical model, along with the lagged convolution method and splined-QWE method. The FHT tests used the 201-point filters derived here as well as the 801-point filters from Anderson (1989). For this particular modeling problem, the 201-point filters are significantly more accurate than Anderson's 801-point filters. Due to the error introduced by the spline interpolation, the lagged-convolutions are significantly less accurate than the direct FHT approach. However, where the ranges are close to the logarithmically spaced filter points, the accuracy of the lagged convolution decreases toward the normal digital filter accuracy. For both the 201- and 801-point filters, the relative error in the lagged convolution at long offsets exceeds 1%, while at shorter ranges the relative error is much lower.

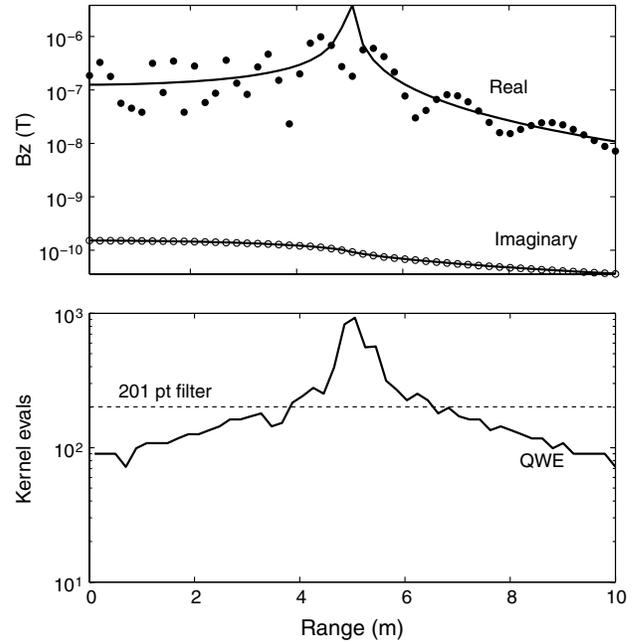


Figure 9. Comparison of the FHT and QWE methods for evaluating the large-loop transform. The top panel shows the real and imaginary responses for a 201-point digital filter (dots) and the QWE method (lines). The singular behavior observed in the real QWE response occurs at the range corresponding to the transmitter coil radius (5 m). The bottom panel shows the number of kernel evaluations required by each method.

For comparison purposes, the splined-QWE method was computed with a relative tolerance of 10^{-2} and a λ density of 40 points per decade for the spline interpolation. The splined-QWE stays below the requested error level over the entire range. This illustrates an advantage of the QWE approach, where the error level can be controlled by adjusting the λ density used for the spline interpolation and the relative error tolerance used for the QWE iterations. Conversely, the spline knot density of the lagged convolution is fixed by the digital filter spacing and therefore cannot be adjusted (unless starting over using another filter, if available).

Now, consider the speed of each approach. Table 1 shows the run-times obtained by each method as a function of the number of offsets. The basic QWE method used a relative tolerance of 10^{-6} while the splined QWE method used 10^{-2} . This was done so that the basic QWE method had errors comparable to the normal FHT, while the splined-QWE method gave errors similar to the lagged convolutions (as shown in Figure 10). Except for the case of only a single offset, the splined-QWE is always faster than the basic QWE method, with the speed-up factor increasing with the number of offsets. However, when more than a few offset are considered, the lagged convolution using

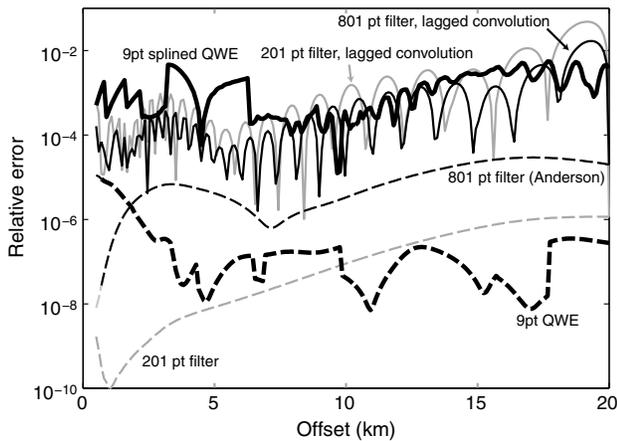


Figure 10. Comparison of the lagged convolution FHT and splined-QWE approaches. The relative error is shown as a function of offset for the normal 201 digital filter, Anderson's 801 point filter, the QWE method, and the corresponding lagged convolution FHTs and the splined-QWE method.

the 201-point filter is much faster than the splined-QWE. For the 21 offset example, the lagged convolution using the 201-point filter is 16 times faster than the normal FHT and almost a factor of two faster than the splined-QWE. The efficiency of the lagged convolution is more dramatic when 321 offsets are considered, with the run-time increasing by less than a factor of two despite the number of offsets increasing by a factor of 15. For this case, the lagged convolution with the 201-point filter is about four times faster than the splined-QWE method. Table 1 shows additional computations for when the number of model layers was increased to 100, which requires significantly more effort for evaluating the kernel functions. Again, the lagged convolution is much faster than the QWE method.

CONCLUSIONS

While the FHT has been the method of choice for the past few decades due to its speed and accuracy, the QWE method presented here is a useful alternative. A significant benefit of the QWE method is the error control available through the adjustable quadrature order and the extrapolation tolerances. This can be used to easily tune the QWE method to give accurate responses for particular model problems and data parameters. Conversely, a given digital filter used for the FHT method provides only a single answer with an unknown accuracy that can vary greatly depending on the specific model and data parameters being considered, although in many cases the FHT digital filters are accurate enough for most practical needs.

For the three EM integral transforms considered here, the QWE method was consistently as fast as or faster than the normal digital filter method used for the FHT. This represents an advancement over a previously proposed quadrature method that uses a continued-fraction expansion and was shown to be significantly slower than the FHT. However, when the integrals are evaluated at a large number of offsets, the lagged convolution variant of the FHT significantly outperforms the QWE method, even when the QWE is accelerated using spline interpolation.

These results suggest the following: When highly accurate solutions are required and speed is not the main concern, the QWE method is the clear choice because the quadrature order and the extrapolation tolerances can be increased until the desired solution accuracy is obtained. The choice is less clear when computational speed is paramount. For a single offset, the QWE method will generally will be faster than the FHT, at least for the types of

Table 1. Comparison of the run times for the normal and lagged convolution FHT methods with the normal and splined QWE methods. Bold numbers show the fastest time for a given number of layers and offsets.

# Layers	# Offsets	QWE		FHT: 201-pt filter		FHT: 801-pt filter	
		Normal	Splined	Normal	Lagged	Normal	Lagged
5	1	0.03	0.06	0.04	0.04	0.15	0.15
5	5	0.14	0.06	0.19	0.05	0.75	0.06
5	21	0.62	0.08	0.80	0.05	3.20	0.17
5	81	2.39	0.15	3.06	0.06	12.15	0.18
5	321	9.45	0.36	12.14	0.09	48.55	0.21
100	21	0.95	0.11	1.24	0.08	4.88	0.26
100	81	3.67	0.18	4.80	0.08	18.76	0.27
100	321	14.49	0.38	19.29	0.11	74.27	0.30

problems considered here. When a large number of offsets are required and they are all at the same depth and for the same source position and orientation, the lagged convolution variant of the FHT is probably the best choice, with the caveat that the solution will have a degraded accuracy due to the spline interpolation. Even this decreased accuracy is likely to be sufficient for most practical applications; however, it should be verified independently (for example, by the QWE method) when encountering a new problem or new range of model parameters. For divergent integrals such as those encountered in the large-loop problem, the QWE method can provide an accurate answer when the FHT method fails.

ACKNOWLEDGMENTS

This work was supported by the Seafloor Electromagnetic Methods Consortium at the Scripps Institution of Oceanography. Eric Kazlauskas and Colin G. Farquharson are thanked for their helpful reviews.

REFERENCES

- Anderson, W. L., 1979, Numerical integration of related Hankel transforms of orders 0 and 1 by adaptive digital filtering: *Geophysics*, **44**, 1287–1305, doi: [10.1190/1.1441007](https://doi.org/10.1190/1.1441007).
- Anderson, W. L., 1982, Fast Hankel-transforms using related and lagged convolutions: *ACM Transactions on Mathematical Software*, **8**, no. 4, 344–368, doi: [10.1145/356012.356014](https://doi.org/10.1145/356012.356014).
- Anderson, W. L., 1984, On: “Numerical integration of related Hankel transforms by quadrature and continued fraction expansion” by A. D. Chave: *Geophysics*, **49**, 1811–1812.
- Anderson, W. L., 1989, A hybrid fast Hankel transform algorithm for electromagnetic modeling: *Geophysics*, **54**, 263–266, doi: [10.1190/1.1442650](https://doi.org/10.1190/1.1442650).
- Chave, A. D., 1983, Numerical integration of related Hankel transforms by quadrature and continued fraction expansion: *Geophysics*, **48**, 1671–1686, doi: [10.1190/1.1441448](https://doi.org/10.1190/1.1441448).
- Chave, A. D., 1984, Reply to discussion of “Numerical integration of related Hankel-transforms by quadrature and continued fraction expansion”, by Chave, A. D. (GEO-48-12-1671-1686): *Geophysics*, **49**, 1813.
- Chen, J., G. M. Hoversten, D. Vasco, Y. Rubin, and Z. Hou, 2007, A Bayesian model for gas saturation estimation using marine seismic AVA and CSEM data: *Geophysics*, **72**, no. 2, WA85–WA95, doi: [10.1190/1.2435082](https://doi.org/10.1190/1.2435082).
- Constable, S., 2010, Ten years of marine CSEM for hydrocarbon exploration: *Geophysics*, **75**, no. 5, 75A67–75A81, doi: [10.1190/1.3483451](https://doi.org/10.1190/1.3483451).
- Ghosh, D., 1971, The application of linear filter theory to the direct interpretation of geoelectrical resistivity sounding measurements: *Geophysical Prospecting*, **19**, 192–217, doi: [10.1111/gpr.1971.19.issue-2](https://doi.org/10.1111/gpr.1971.19.issue-2).
- Guptasarma, D., and B. Singh, 1997, New digital linear filters for Hankel J(0) and J(1) transforms: *Geophysical Prospecting*, **45**, 745–762, doi: [10.1046/j.1365-2478.1997.500292.x](https://doi.org/10.1046/j.1365-2478.1997.500292.x).
- Haines, G., and A. G. Jones, 1988, Logarithmic Fourier transformation: *Geophysical Journal International*, **92**, 171–178, doi: [10.1111/gji.1988.92.issue-1](https://doi.org/10.1111/gji.1988.92.issue-1).
- Hamming, R. W., 1986, *Numerical methods for scientists and engineers*: Dover Publications.
- Hänggi, P., F. Rosel, and D. Trautmann, 1978, Continued fraction expansions in scattering theory and statistical nonequilibrium mechanics: *Zeitschrift Fur Naturforschung*, **33**, 402–417.
- Johansen, H. K., and K. Sorensen, 1979, Fast Hankel transforms: *Geophysical Prospecting*, **27**, 876–901, doi: [10.1111/gpr.1979.27.issue-4](https://doi.org/10.1111/gpr.1979.27.issue-4).
- Key, K., 2009, 1D inversion of multicomponent, multifrequency marine CSEM data: Methodology and synthetic studies for resolving thin resistive layers: *Geophysics*, **74**, no. 2, F9–F20, doi: [10.1190/1.3058434](https://doi.org/10.1190/1.3058434).
- Kong, F. N., 2007, Hankel transform filters for dipole antenna radiation in a conductive medium: *Geophysical Prospecting*, **55**, 83–89, doi: [10.1111/gpr.2007.55.issue-1](https://doi.org/10.1111/gpr.2007.55.issue-1).
- Li, Y., and K. Key, 2007, 2D marine controlled-source electromagnetic modeling: Part 1 — An adaptive finite element algorithm: *Geophysics*, **72**, no. 2, WA51–WA62, doi: [10.1190/1.2432262](https://doi.org/10.1190/1.2432262).
- Longman, I. M., 1956, Note on a method for computing infinite integrals of oscillatory functions: *Mathematical Proceedings of the Cambridge Philosophical Society*, **52**, no. 4, 764–768, doi: [10.1017/S0305000410003187X](https://doi.org/10.1017/S0305000410003187X).
- Lucas, S. K., and H. A. Stone, 1995, Evaluating infinite integrals involving Bessel functions of arbitrary order: *Journal of Computational and Applied Mathematics*, **64**, no. 3, 217–231, doi: [10.1016/0377-0427\(95\)00142-5](https://doi.org/10.1016/0377-0427(95)00142-5).
- Michalski, K. A., 1998, Extrapolation methods for Sommerfeld integral tails: *Antennas and Propagation*, **140**, no. 10, 5–1418, doi: [10.1109/8.725271](https://doi.org/10.1109/8.725271).
- Newman, G., G. Hohmann, and W. Anderson, 1986, Transient electromagnetic response of a three-dimensional body in a layered earth: *Geophysics*, **51**, 1608–1627, doi: [10.1190/1.1442212](https://doi.org/10.1190/1.1442212).
- Newman, G. A., and D. L. Alumbaugh, 1995, Frequency-domain modelling of airborne electromagnetic responses using staggered finite differences: *Geophysical Prospecting*, **43**, 1021–1042, doi: [10.1111/gpr.1995.43.issue-8](https://doi.org/10.1111/gpr.1995.43.issue-8).
- Patterson, T. N. L., 1968, The optimum addition of points to quadrature formulae: *Mathematics of Computation*, **22**, no. 104, 847–856, doi: [10.1090/S0025-5718-68-99866-9](https://doi.org/10.1090/S0025-5718-68-99866-9).
- Raiche, A., 1999, A flow-through Hankel transform technique for rapid, accurate Green’s function computation: *Radio Science*, **34**, no. 2, 549–555, doi: [10.1029/1998RS900037](https://doi.org/10.1029/1998RS900037).
- Shanks, D., 1955, Nonlinear transformations of divergent and slowly convergent sequences: *Journal of Mathematical Physics*, **34**, 1–42.
- Sorensen, K., and N. B. Christensen, 1994, The fields from a finite electrical dipole — A new computational approach: *Geophysics*, **59**, 864–880, doi: [10.1190/1.1443646](https://doi.org/10.1190/1.1443646).
- Trefethen, L. N., 2000, *Spectral methods in MATLAB*: SIAM, volume **10** of *Software, Environments, and Tools*.
- Ward, S. H., and G. W. Hohmann, 1987, Electromagnetic theory for geophysical applications, in *Electromagnetic Methods in Applied Geophysics*: SEG, 1–312.
- Weniger, E. J., 1989, Nonlinear sequence transformations for the acceleration of convergence and the summation of divergent series: *Computer Physics Reports*, **10**, 189–371, doi: [10.1016/0167-7977\(89\)90011-7](https://doi.org/10.1016/0167-7977(89)90011-7).
- Wynn, P., 1956, On a device for computing the $e_m(S_n)$ transformation: *Mathematical Tables and Other Aids to Computation*, **10**, no. 54, 91–96.