Fast Fourier transform of electromagnetic data for computationally expensive kernels

Werthmüller, Dieter; Mulder, Wim A; Slob, Evert C

DOI

[10.1093/gji/ggab171](https://doi.org/10.1093/gji/ggab171)

Publication date

2021

Document Version

Final published version

Published in

Geophysical Journal International

Citation (APA)

Werthmüller, D., Mulder, W. A., & Slob, E. C. (2021). Fast Fourier transform of electromagnetic data for computationally expensive kernels. *Geophysical Journal International*, 226(2), 1336-1347. [ggab171]. <https://doi.org/10.1093/gji/ggab171>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Fast Fourier transform of electromagnetic data for computationally expensive kernels

Dieter Werthmüller¹, Wim A. Mulder^{1,2} and Evert C. Slob¹

¹TU Delft, Building 23, Stevinweg 1 / PO-box 5048, 2628 CN Delft, Netherlands. E-mail: dieter@werthmuller.org

²Shell Global Solutions International BV, Grasweg 31, 1031 HW Amsterdam, Netherlands

Accepted 2021 April 20. Received 2021 April 7; in original form 2020 June 30

SUMMARY

3-D controlled-source electromagnetic data are often computed directly in the domain of interest, either in the frequency domain or in the time domain. Computing it in one domain and transforming it via a Fourier transform to the other domain is a viable alternative. It requires the evaluation of many responses in the computational domain if standard Fourier transforms are used. This can make it prohibitively expensive if the kernel is time-consuming as is the case in 3-D electromagnetic modelling. The speed of modelling obtained through such a transform is defined by three key points: solver, method and implementation of the Fourier transform, and gridding. The faster the solver, the faster modelling will be. It is important that the solver is robust over a wide range of values (frequencies or times). The method should require as few kernel evaluations as possible while remaining robust. As the frequency and time ranges span many orders of magnitude, the required values are ideally equally spaced on a logarithmic scale. The proposed fast method uses either the digital linear filter method or the logarithmic fast Fourier transform together with a careful selection of evaluation points and interpolation. In frequency-to-time domain tests this methodology requires typically 15–20 frequencies to cover a wide range of offsets. The gridding should be frequency- or time-dependent, which is accomplished by making it a function of skin depth. Optimizing for the least number of required cells should be combined with optimizing for computational speed. Looking carefully at these points resulted in much smaller computation times with speedup factors of ten or more over previous methods. A computation in one domain followed by transformation can therefore be an alternative to computation in the other domain if the required evaluation points and the corresponding grids are carefully chosen.

Key words: Controlled source electromagnetics (CSEM); Fourier analysis; Numerical modelling.

1 INTRODUCTION

The controlled-source electromagnetic (CSEM) method is one of the common non-seismic tools in exploration geophysics, not only in hydrocarbon exploration (Constable 2010), but also in the search for sulfides (Gehrmann *et al.* 2019), water (Pedersen *et al.* 2005), geothermal sources (Girard *et al.* 2015) or for geological purposes (Johansen *et al.* 2019). While current sources with a few frequencies are used in the deep marine environment, transient measurements are more common in the shallow marine environment and on land (e.g. Andréis & MacGregor 2007; Avdeeva *et al.* 2007; Ziolkowski *et al.* 2007). One of the main reasons is the dominance of the air-wave in shallow marine and terrestrial measurements, which can be better separated in the time domain. CSEM is usually divided into frequency- and time-domain methods, depending on whether the source signal is a continuous waveform, such as a sine, or a

finite waveform, such as a pseudo-random binary sequence (PRBS). A numerical comparison of the two methods is given by Connell & Key (2013). Acquired CSEM data are subsequently often analysed (modelled and inverted for) in their respective domain, either the frequency or time domain. Modellers of layered media usually exploit the horizontal shift-invariance by computing the responses in the wavenumber-frequency domain followed by a 2-D inverse spatial Fourier transform, also called Hankel transform, to the space-frequency domain, and a regular inverse Fourier transform if time-domain data are required (e.g. Hunziker *et al.* 2015). CSEM codes for arbitrary 2-D and 3-D computations, on the other hand, often compute their responses directly in the required domain, either frequency or time.

Electromagnetic methods in geophysics span a wide range of model scenarios and acquisition layouts, each with its own modelling-related implications. A well-known example is the

mentioned land or shallow marine case versus the deep marine case, where in the former case a dominating airwave has to be accurately modelled, whereas it can be completely ignored in the latter case. Recent areas of particular interest in time-domain modelling that pose numerical challenges are, for instance, simulating the fields through steel-cased wells (Heagy & Oldenburg 2019) or the effects of induced polarization (Kang *et al.* 2020). The former is challenging because of the high conductivity contrasts requiring very detailed meshing, the latter is challenging for time-domain modelling because the models are frequency dependent.

Under certain conditions it can be a viable alternative to model time-domain data with a frequency-domain code, as shown by Mulder *et al.* (2008). They compared the computational complexity of different time-domain methods with the computational complexity of frequency-domain computation followed by a Fourier transform: explicit time-stepping schemes (Du Fort-Frankel method as used in, e.g. Commer & Newman 2004; Maaø 2007) and implicit schemes (e.g. Haber *et al.* 2004; Um *et al.* 2010) have a complexity of $O(n^4)$, matrix exponentials and Lanczos reduction schemes (e.g. Druskin & Knizhnerman 1994) have a complexity of $O(n^4 \sqrt{\log n})$, and Fourier transforms of frequency-domain solutions have a complexity $O(n_f n^3)$; n is the number of gridpoints in each direction, and n_f the number of frequencies. Their conclusion was that the Fourier transform method can be favourable if the number of required frequencies is small relative to the number of gridpoints in each coordinate. The conditions for fast computation of time-domain data with a frequency-domain code are: a sufficiently powerful solver, appropriate frequency selection and interpolation, and an automated gridding, for which they used the multifrequency CSEM approach presented by Plessix *et al.* (2007). We build upon these results but improve the run time from hours to minutes. The main reasons for this significant speed-up are the further reduction of required frequencies by introducing lower and upper thresholds of numerically important frequencies, an adaptive, frequency-dependent gridding scheme that minimizes the required cells in each dimension, and a logarithmic Fourier transform such as digital linear filters (DLF, Ghosh 1971) or the logarithmic fast Fourier transform (FFTLog, Hamilton 2000) to go from the frequency to the time domain. The latter makes it also possible to only use the imaginary part of the frequency-domain response, which has advantages when it comes to interpolation.

In the next section, we briefly review the methodology as introduced by Plessix *et al.* (2007) and Mulder *et al.* (2008) and highlight their advantages and shortcomings. This is followed by an outline of the methodology, our changes to the Fourier transform and the adaptive gridding. Finally, we demonstrate the efficiency of the approach with some numerical results. These comprise in a homogeneous space, a layered model, a model which includes dispersive media (induced polarization), and a 3-D model.

2 MOTIVATION

Being able to model CSEM data both in the frequency domain and in the time domain can be desirable, as both domains have advantages and disadvantages. One way to achieve this is to implement Maxwell's equations in both domains as is done, for instance, in SimPEG (Cockett *et al.* 2015) or in custEM (Rochlitz *et al.* 2021). Another approach is to have Maxwell's equations only implemented in one domain, and use Fourier transforms to go to the other. However, this approach can be costly, as many frequencies over a wide range are required to go from the frequency domain to the time

domain, or many times over a wide range for the opposite direction. We present a methodology which significantly reduces the amount and range of the required frequencies through a combination of extrapolation, interpolation, and setting the responses for high frequencies to zero. This translates into a significant reduction in computation time. The number of required frequencies is not the only important point for reducing computation time. Another important aspect is the computation grid, as the required grids for low frequencies (in our case in the order of 0.001 Hz) and high frequencies (around 100 Hz) are hugely different. Low frequencies can be computed on a coarser grid, but they require a much larger domain in order to avoid boundary effects. High frequencies, on the other hand, require denser gridding, but they are much more limited in reach. An adaptive gridding scheme is therefore required, which is naturally based on the skin depth, the distance after which the amplitude of the electromagnetic field has decayed by $1/e \approx 37\%$. The skin depth δ is a function of conductivity and frequency, and for the diffusive approximation of Maxwell's equations in an isotropic, homogeneous medium is given by (e.g. Ward & Hohmann 1988)

$$\delta = \sqrt{\frac{2}{\omega \mu \sigma}} \approx 503.3 / \sqrt{f \sigma}, \quad (1)$$

where σ is conductivity (S/m), $\omega = 2\pi f$ is angular frequency of frequency f (Hz) and μ is magnetic permeability (H/m). The approximation is obtained by using the free-space value of magnetic permeability, $\mu_0 = 4\pi \times 10^{-7}$ H/m.

We build our approach upon Plessix *et al.* (2007), who presented such an adaptive gridding for multifrequency (and multisource) CSEM modelling. They define the minimum cell width $\Delta_{\min}$ as a fraction of the minimum skin depth $\delta_{\min}$, where $\delta_{\min}$ should be two to three times bigger than $\Delta_{\min}$. The cells have to be smallest around the source; in the marine case, the minimum skin depth is therefore defined by the conductivity of seawater. However, this can yield quite large cells for low frequencies, so special care has to be taken around the source by defining a maximum allowed $\Delta_{\min}$. The grid dimension, on the other hand, is defined as a function of skin depth for the average conductivity of the background, δ_{ave} . They use four times δ_{ave} for the x -, y - and downward z -directions, and a fixed 50 km for the upward z -direction to account for the airwave. To reduce the number of cells, it is desirable to introduce stretching, at least in the buffer zone outside of the area where source and receivers are located.

Mulder *et al.* (2008) provide a computational complexity analysis of various methods to model transient electromagnetic responses directly in the time domain, and compare it to the computation of transient EM responses in the frequency domain with a subsequent Fourier transform. They conclude the review by stating “*Although it remains to be seen which of the four methods requires the least computer time for a given accuracy, the frequency-domain approach appears to be attractive.*” Their approach is to minimize the computation time by having, in addition to the just introduced adaptive gridding, an adaptive frequency selection scheme. This scheme starts with computing the responses for a set of just five frequencies, regularly sampled on a log-scale, from minimum to maximum required frequency. All the other frequencies are interpolated with a shape-preserving piecewise-cubic Hermite interpolation (PCHIP, Fritsch & Carlson 1980). Testing the stability of the obtained response by removing a single frequency-value at a time their scheme decides if more frequencies in-between the already computed ones are required. In this way frequencies are only added if required, hence if certain criteria of response stability are not met. While this

method is good and effective for a single offset, it loses all its advantages if one tries to compute different offsets within one computation, as each offset requires a different set of adaptive frequencies. Additionally, it hampers the parallelization over frequencies.

We present improvements to both the adaptive gridding and the transform from frequency domain to time domain using the finite-integration technique (Weiland 1977), which makes time-domain modelling with a frequency-domain code viable in comparison with time-domain codes. It is important to note that our ideas can be used with any solver and is not bound to one or another frequency- or time-domain solver. The ideas for the frequency selection can be applied to time selection, and time-dependent, adaptive gridding exists as well (e.g. Commer & Newman 2006). Our recommendations for speeding up the Fourier transform for expensive kernels are independent of spatial complexity, as diffusive EM responses are smooth functions of both frequency and time. However, the intrinsic requirement is an accurate solver over a sufficient wide range of frequencies or times.

3 METHODOLOGY

The requirement for any transform is a robust solver of sufficient accuracy over a wide range of frequencies. We use for the numerical computations the open-source (Apache License 2.0) codes *empymod* (Werthmüller 2017) and *emg3d* (Werthmüller *et al.* 2019b). The former computes semi-analytical responses of layered models, the latter is a 3-D multigrid solver based on Mulder (2006), which can be used as a pre-conditioner for Krylov subspace solvers or as a solver on its own. The multigrid approach works fine for the diffusive approximation of Maxwell's equations, which assumes that $\omega\epsilon \ll \sigma$, where ϵ is electric permittivity (F/m). The remaining system to solve in the frequency domain is then given by the second-order differential equation of the electric field,

$$i\omega\sigma\mathbf{E} + \nabla \times \mu^{-1} \nabla \times \mathbf{E} = -i\omega\mathbf{J}_s, \quad (2)$$

where $\mathbf{E}$ is the electric field (V/m) and $\mathbf{J}_s$ the current source (A/m²); time dependence is $\exp(i\omega t)$. The standard multigrid approach fails for severe stretching or strong anisotropy, for which known improvements such as line-relaxation and semicoarsening (Jönsthövel *et al.* 2006) are implemented, with a non-standard Cholesky decomposition to speed up the computation (Mulder *et al.* 2008). One of the big advantages of the multigrid method is that it scales linearly (optimal) with the grid size in both CPU and RAM usage (Mulder 2020). This makes it feasible to run even big models on standard computers, without the need for big clusters. All examples in this article are run on a laptop with an i7-6600U CPU @ 2.6 GHz ($\times 4$) and 16 GB of memory, using Ubuntu 20.04 and Python 3.8. They were run using a single thread, but parallelization over frequencies is straight forward, as the computations are independent.

Note that while we use these two algorithms for the numerical examples, any solver that provides sufficiently accurate frequency-domain responses could be used.

3.1 Frequency selection

An important factor in terms of speed and accuracy for time-domain responses obtained from frequency-domain computations is the selection of the required frequencies. The fewer frequencies required, the quicker we obtain the time-domain result. This is true in general but applies in particular to expensive kernels such as is the case in 3-D modelling, where the computation of the response of

a single frequency takes much longer than the transformation itself. The same applies for the reverse operation, the fewer times required, the quicker we obtain frequency-domain results. Even though we limit our analysis here to frequency-to-time transformations, most of the arguments are reversible and applicable to time-to-frequency transformations. We decided to use a regular spacing of frequencies on a log-scale, rather than an adaptive scheme. This approach is favourable if a wide range of offsets is needed, as the required frequencies change with offset and an adaptive frequency selection is therefore often tailored to a single offset or a limited range of offsets. Also, it allows for straightforward parallelization over frequencies, which is not completely possible with an adaptive scheme.

For the actual transform we use either the digital linear filter (DLF) method or the logarithmic fast Fourier transform (FFTLog). The DLF method was introduced to geophysics by Ghosh (1971), and is arguably the most common method in EM geophysics for both its simplicity and its speed. It is implemented for the Hankel and Fourier transforms in most EM modelling codes (e.g. Key 2009). A simple tool to design digital linear filters was recently presented by Werthmüller *et al.* (2019a), together with a comprehensive overview of the history and development of DLF in geophysics. FFTLog, introduced by Hamilton (2000), is another transform algorithm which proved to be powerful for the frequency-to-time domain transformation of EM responses (e.g. Werthmüller *et al.* 2014). In our tests they are both about equal in speed and accuracy. DLF requires a wider range and many more frequencies than the FFTLog. Both methods share some important characteristics in comparison with the standard FFT: the required input frequencies are equally spaced on a logarithmic scale (natural logarithm in the case of the DLF and decimal logarithm in the case of the FFTLog), and they only require either the real or the imaginary part of the frequency-domain response. We can take advantage of that by using only the imaginary part of the frequency-domain response. The imaginary part goes to zero when the frequency goes either to zero or to infinity, with the advantage that knowing the endpoints makes it possible to convert the extrapolation of missing frequencies into an interpolation.

Any Fourier transform needs a certain range of frequencies to obtain time-domain results. Our approach is not to minimize the required frequencies for the transform, but to minimize the number of frequencies for which we actually have to compute responses. We do this in three zones: above an upper threshold $f_{\max}$, we set the frequency-domain response to zero. Below a lower threshold $f_{\min}$, we extrapolate the frequency-domain response. As we know that the imaginary response goes towards zero for zero frequency, we can turn it into an interpolation, for which we use PCHIP. And in-between the thresholds, we use cubic spline interpolation in order to only have to actually compute a few distinct frequencies per decade (n_{dec}^f). Using this approach, we found that computing the responses for 15–25 frequencies is usually good enough to obtain time-domain responses for a large range of offsets; even though the actual transform might need up to hundreds of frequencies.

The required frequencies depend naturally on the conductivities and on the chosen acquisition type and layout. A trial-and-error approach with a 3-D code is very time-intensive. However, a simplified, layered model for the required survey setup and a fast 1-D modeller makes it possible to estimate these parameters easily. An example of this is shown in Fig. 1. It shows the responses for a marine scenario with 1 km water depth of resistivity $\rho = 0.3 \Omega \text{ m}$ ($\rho = \sigma^{-1}$) and a 100-m-thick target of $100 \Omega \text{ m}$ at 1 km below the seafloor in a background of $1 \Omega \text{ m}$. The source is 50 m above the

Figure 1. Example frequency selection for a simple layered model (the shown model parameters are described in the text). This example shows the impulse response at an offset of 5 km, for which it uses FFTLog from 0.001 to 10 Hz with five frequencies per decade.

seafloor, the receiver is on the seafloor, and the response is the inline x -directed E -field. The left-hand subplot shows the imaginary part in the frequency domain and the right-hand subplot the corresponding impulse response in the time domain. The red lines are the semi-analytical responses. The blue circles indicate the actually computed responses and the black dots the frequencies which are interpolated or set to zero. The resulting time-domain response has a relative error of $<1\%$ everywhere except for very early times.

Fig. 2(a) shows exactly the same on a linear scale for the time and amplitude axes. It can be seen that with the chosen, computed frequencies (blue dots) the time-domain response starts to divert above about 100 s and below 0.3 s. It also shows the oscillating high-frequency part, which is hard to interpolate and we therefore set it to zero. Fig. 2(b) is the same as Fig. 2(a), but transformed with the DLF method applying the 81-point sine-cosine filter from Key (2009). The same frequencies were computed as in the FFTLog case and the missing ones interpolated. The error of the corresponding time-domain response is comparable, so either FFTLog or DLF can be used. In the above example we used 20 frequencies, but that many would not be required for the shown offset of 5 km, a few of the lower and higher frequencies could be left out. With this frequency-selection, however, we can model a wide range of offsets. This is shown in Fig. 3, where the same kernels were used to yield the responses at offsets $r = 1.5, 3, 6$ and 12 km. We only need to compute 50 % of the required frequencies in the case of the FFTLog, and only 14 % in the case of the DLF.

The shortest offset defines the highest required frequency, and the largest offset the lowest required frequency. So the more one can restrict the necessary offset range, the fewer frequencies are needed. Another important factor is how to interpolate and extrapolate from the computed frequencies to the frequencies required for the Fourier transform. For the FFTLog only extrapolation for higher and lower frequencies is required. The EM response becomes highly oscillatory for high frequencies, which makes it very hard to extrapolate the response to frequencies $f > f_{\max}$. However, if $f_{\max}$ is chosen judiciously, the importance of higher frequencies for the Fourier transform can be neglected and we can set those responses to zero. The extrapolation of frequencies $f < f_{\min}$ can be changed to an interpolation by assuming a zero imaginary response at zero frequency, and we then use PCHIP to interpolate the missing frequencies (as

we work on a logarithmic scale we cannot choose 0 Hz nor 0 V/m, but instead take 10^{-100} Hz and 10^{-100} V/m). In the case of the DLF method we also have to interpolate in between the computed frequencies, for which we found it better to use a cubic spline. As can be seen by comparing Fig. 2(a) with Fig. 2(b), using the FFTLog or the DLF with the same actually computed frequencies results in very similar responses.

The changes to the Fourier transform in comparison with Mulder *et al.* (2008) can be summarized in three points: (1) regular log-scale spacing for the frequency selection; (2) DLF or FFTLog instead of FFT and (3) using only the imaginary part of the frequency-domain response. The actual speed of the transform is unimportant, as the computation of the frequency-domain responses takes much longer than the transform itself. What matters is solely how many frequencies are required by it to achieve the desired precision, and how long it takes to compute the responses for these frequencies.

3.2 Gridding

The proposed Fourier transform requires a robust solver that can compute accurate results over a wide range of frequencies (times) to obtain time-domain (frequency-domain) responses. In our case the solver uses a regular, stretched grid and computes the electric fields in the frequency domain. This setup is the target of the following gridding recommendations. Whilst they will look differently for other mesh types or a time-domain code, some conclusions will still hold.

The computation grid consists of a core or survey domain D_s that should contain all source and receiver positions. The survey domain usually has no or a very small cell stretching factor α_s . The minimum cell width is defined as

$$\Delta_{\min} = \delta(f, \sigma_{\text{src}}) / n_\delta, \quad (3)$$

where σ_{src} is the conductivity of the media in which the source resides, and n_δ is a positive number that defines how many cells there should be per skin depth. The actual computational domain D_c is usually much bigger than D_s in order to avoid artefacts from the boundary condition. It can also have a much higher stretching factor α_c . In the presented examples we have chosen D_c such that the distance for the signal diffusing from the source to the boundary

Figure 2. (a) Same as Fig. 1, but on a logarithmic scale. (b) Using DLF instead of FFTLog, but computing the same frequency range as for the FFTLog. The resulting time-domain response has comparable accuracy.

Figure 3. Normalized (a) frequency- and (b) time-domain responses for offsets $r = 1.5, 3, 6$ and 12 km using the parameters defined in Fig. 2(a). The coloured symbols are the actually computed responses, the black dots are the responses which are set to zero (high frequencies) or interpolated (low frequencies). The grey curves are the analytical responses.

and back to the receiver closest to the boundary is at least two wavelengths, after which the amplitude of the signal is reduced to a millionth of its initial strength. The wavelength λ to compute D_c is given by

$$\lambda = 2\pi\delta(f, \sigma_{\text{ave}}) \approx 3162/\sqrt{f\sigma_{\text{ave}}}, \quad (4)$$

where σ_{ave} is the average conductivity, which can vary for different directions. However, the skin-depth approach fails for air, in which the EM field propagates as a wave at the speed of light. A largest computational domain is therefore enforced, defining the maximum distance from the source to the boundary; this distance is by default set to 100 km, but this can be reduced in the marine case with increasing water depth. This also applies to the horizontal dimensions, not only to the upward z -direction, and it equally applies to very resistive basements, even in deep water. One way to circumvent this difficulty is the use of a primary-secondary formulation, where the primary field, including the air wave, is computed with a semi-analytical code for layered media. We do not consider this approach here. Note that grid stretching for complex-valued diffusion fields is essentially what is done for wave fields with the perfectly

matched layer (PML). PML is an absorbing layer to avoid scattering from the boundary by letting the field decay to zero, which is achieved by introducing a complex factor that causes damping. As electromagnetic diffusion fields are damped fields by themselves, it suffices to stretch the grid.

In summary, the adaptive gridding takes $f, D_s, \sigma_{\text{src}}, \sigma_{\text{ave}}, n_\delta$ and ranges for α_s, α_c , where we usually fix $\alpha_s = 1$ or keep it at least below 1.05, and let α_c be anything between $[1, 1.5]$. The minimum cell width Δ_{min} can further be restricted by a user-defined range. Given these inputs, the adaptive gridding will search for the smallest possible number of cells which fulfils these criteria. The implemented multigrid method puts some constraints on the number of cells, of which the adaptive gridding takes care (the number of cells have to be powers of two multiplied by a low prime, for example $\{2, 3, 5\} \cdot 2^n$).

The main difference with Mulder *et al.* (2008) is that their adaptive gridding searches for the optimal stretching factor α fulfilling certain criteria, for a fixed number of cells. Our adaptive gridding, on the other hand, searches for the smallest number of cells that still fulfils the given criteria. The number of cells becomes therefore

Table 1. Run times per frequency for the homogeneous space example, with the corresponding number of cells and minimum cell width as well as the stretching factor in the computation domain; $\alpha_s = 1$ everywhere.

Frequency (Hz)	Time (s)	$n_x \times n_y \times n_z$	$\Delta_{\min}$ (m)	α_c
20.0	4	$80 \times 24 \times 24$	20	1.26
12.6	8	$96 \times 32 \times 32$	20	1.17
7.98	8	$96 \times 32 \times 32$	20	1.20
5.03	8	$96 \times 32 \times 32$	20	1.23
3.18	7	$80 \times 32 \times 32$	24	1.23
2.00	7	$80 \times 32 \times 32$	30	1.21
1.26	5	$64 \times 32 \times 32$	37	1.21
0.798	5	$64 \times 32 \times 32$	40	1.23
0.503	5	$64 \times 32 \times 32$	40	1.26
0.318	5	$64 \times 32 \times 32$	40	1.28
0.200	8	$64 \times 40 \times 40$	40	1.27
0.126	8	$64 \times 40 \times 40$	40	1.30
0.0798	10	$80 \times 40 \times 40$	40	1.26
0.0503	11	$80 \times 40 \times 40$	40	1.28

also a function of frequency. It is important to note that this is our implementation of an adaptive grid, but there are certainly other possibilities. The relevant point for fast computations is that the adaptive gridding tries to minimize the number of required cells. This is generally best done in a frequency- and conductivity-dependent manner. To go from the model grid to the computational grid, we use the volume-averaging technique on logarithmic resistivities, as used in Plessix *et al.* (2007). While this technique ensures that the total resistivity in the subsurface remains the same, it does not consider effective-medium theory (Davydycheva *et al.* 2003), for instance, the apparent anisotropy from a stack of finely layered formations of varying resistivity.

4 NUMERICAL EXAMPLES

4.1 Homogeneous space

The first example is the inline electric field from a source at the origin measured by an inline receiver with an offset of 900 m in a homogeneous space of $1 \Omega \text{ m}$. We chose this simple example to compare it with the analytical solution and with previously published results. We used the following values to define the required frequencies: $f_{\min} = 0.05 \text{ Hz}$, $f_{\max} = 21 \text{ Hz}$, using FFTLog with five frequencies per decade. This results in 14 frequencies to compute from 0.05 to 20.0 Hz. The complete frequency range for the transform, including the frequencies for which we use interpolation, includes 30 frequencies from 0.0002 to 126.4 Hz. For the adaptive gridding the following inputs were used: $n_\delta = 12$, minimum cell width must be between 20 and 40 m, and $\alpha_s = 1$, $\alpha_c = [1, 1.3]$. This yields grids with cell numbers between 46 080 ($80 \times 24 \times 24$, for 20.0 Hz) and 128 000 ($80 \times 40 \times 40$, for 0.05 Hz) cells. The run times for each frequency, the corresponding number of cells, minimum cell width, and computation domain stretching factor are listed in Table 1. The total run time to compute this model was less than 2 min.

Fig. 4(a) shows the frequency-domain result, where the blue dots are the computed responses and the black dots correspond to the interpolated values or the values set to zero. Most of the computed values stay below a relative error of 1 %, our chosen adaptive gridding only starts to generate considerable errors at higher frequencies. Fig. 4(b) shows the corresponding time-domain result, where the dashed black line is the result obtained by transformation

of the frequency-domain response, on top of the red line which is the analytical result. The relative error is mostly below 1 %, except for early times. However, for practical reasons that is more than enough. Fig. 5 shows the same on a logarithmic scale, with times up to 10 s. It clearly shows that if later times are required, we would need to adjust our Fourier transform parameters. Note that for the gridding we chose $n_\delta = 12$, which is very dense. This was necessary because we are relatively close to the source. If the offsets of interest are larger this factor can be lowered considerably; 3–4 is often enough.

This model corresponds to the one presented in Table 1 and in figs 3 and 4 of Mulder *et al.* (2008). The response here appears to be more accurate, their reported peak-error is roughly 1 %, whereas we are below 0.1 % at the peak (there are no error-plots presented, so visual inspection is all we have). However, the difference in run time is dramatic. Summing the run times for the different frequencies of the original figure comes to a total computation time of 3 hr 47 min 12 s; 0.01 Hz was the slowest run with 31 min 19 s, and 2.37 Hz was the fastest run with 2 min 54 s. Our example, on the other hand, took less than two minutes in total, where the individual frequencies took between 4 and 11 s to run.

This massive speed-up has a couple of reasons. Computers have become more powerful in the last 12 yr, and the codes were run on different computers. A quick test with the old scripts on our test machine shows that it would roughly run two to three times faster, therefore somewhere between 1 and 2 hr. The more important facts besides different hardware are: (1) we only used 14 frequencies instead of the 26 frequencies between 0.01 and 100 Hz of the original and (2) our adaptive gridding used significantly less cells (f -dependent) in comparison to the fixed 2097 152 cells (128^3) used in the original example. We did not see a significant difference in the speed of the actual codes, where the kernel-algorithm of the two implementations is the same, but in the original example it is implemented in Matlab/C, whereas *emg3d* is written in Python/Numba [Numba is a just-in-time compiler for Python code, Lam *et al.* (2015)].

4.2 1-D model

The second example is a shallow marine, layered model with 200 m of seawater (3 S/m) above a half-space of 1 S/m, an embedded target layer at 2 km depth, 100 m thick, with a conductivity of 0.02 S/m. The source is located 20 m above the seafloor and the receivers are on the seafloor. We chose the frequency range such that we can model offsets from 3 to 7 km, with $f_{\min} = 0.007 \text{ Hz}$ and $f_{\max} = 32 \text{ Hz}$, using FFTLog with five frequencies per decade. This results in computations for 19 frequencies from 0.008 to 31.8 Hz. The complete frequency range for the transform includes 35 frequencies from 2×10^{-5} to 126.4 Hz. For the adaptive gridding, we used a cell width of 100 m in the core domain and stretching outside up to a factor 1.5, where the computation domain extends up to 50 km in each direction. This yielded grids between 204 800 (higher frequencies) and 245 760 (lower frequencies) cells. The run times for each frequency and their corresponding parameters are listed in Table 2.

Fig. 6 shows the result for an offset of 5 km, (a) in the frequency domain and (b) in the time domain. The recovered response with the 3-D code captures the airwave (first peak) and the subsurface (second peak) very accurately. At later times the error starts to increase. We would need to compute a few additional lower frequencies if we want to improve it. In the frequency-domain plot, it can be seen that

Figure 4. (a) Frequency- and (b) time-domain results for the homogeneous space model. The red lines are the analytical solutions, the blue circles are the actually computed responses, the black dots are the interpolated responses and the dashed black line the obtained time-domain response.

Figure 5. Same as in Fig. 4 (b), but on a logarithmic scale. To improve later times we would have to compute lower frequencies; to improve earlier times we would have to compute more frequencies per decade to get a better resolution.

the high frequencies are not computed very accurately, but without too much influence on the time-domain response. These frequencies could be left out if an offset of 5 km is the only objective. However, we also want to retrieve shorter offsets from the same computation, for which these frequencies are required.

Fig. 7 shows the time-domain responses of the same model for offsets of 3, 5 and 7 km, all obtained with the same frequency-domain computations and the same frequencies for the Fourier

Table 2. Run times per frequency for the marine 1-D example, with the corresponding number of cells and minimum cell width as well as the stretching factor in the computation domain; $\alpha_s = 1$ everywhere.

Frequency (Hz)	Time (s)	$n_x \times n_y \times n_z$	$\Delta_{\min}$ (m)	α_c
31.8	16	$128 \times 40 \times 40$	100	1.36
20.0	17	$128 \times 40 \times 40$	100	1.36
12.6	17	$128 \times 40 \times 40$	100	1.36
7.98	17	$128 \times 40 \times 40$	100	1.36
5.03	17	$128 \times 40 \times 40$	100	1.44
3.18	17	$128 \times 40 \times 40$	100	1.48
2.00	17	$128 \times 40 \times 40$	100	1.49
1.26	20	$128 \times 40 \times 48$	100	1.36
0.798	24	$128 \times 40 \times 48$	100	1.36
0.503	27	$128 \times 40 \times 48$	100	1.36
0.318	27	$128 \times 40 \times 48$	100	1.36
0.200	34	$128 \times 40 \times 48$	100	1.38
0.126	34	$128 \times 40 \times 48$	100	1.40
0.0798	37	$128 \times 40 \times 48$	100	1.41
0.0503	37	$128 \times 40 \times 48$	100	1.44
0.0318	44	$128 \times 40 \times 48$	100	1.44
0.0200	44	$128 \times 40 \times 48$	100	1.47
0.0126	37	$128 \times 40 \times 48$	100	1.48
0.00798	43	$128 \times 40 \times 48$	100	1.49

transform. The computation of these frequencies took less than 9 min and it handles any offset between 3 and 7 km. It can be seen that the chosen frequency selection is sufficient for this offset range; again, more low frequencies could be added to improve late-time values.

4.3 Horizontal extent of the computation domain

The skin-depth approach fails for the air layer, as explained in Section 3.2. The reason is that the EM field in the air travels at the

Figure 6. Electric inline response at an offset of 5 km for a shallow marine, layered scenario. (a) Frequency-domain response, where the blue circles denote computed responses and the black dots interpolated responses or responses set to zero. (b) Time-domain response.

Figure 7. Time-domain responses for offsets of 3, 5 and 7 km for the same model as shown in Fig. 6.

speed of light as a wave, and its amplitude is only reduced through geometrical spreading. On land and in shallow marine scenarios one has therefore to include a sufficiently large computational domain. The default in our scheme is 100 km. The important point is that this does not only apply to the upward z -direction, but also to the horizontal directions, as the airwave bounces back horizontally and would continuously emit energy into the subsurface if the boundaries are not chosen far enough away from the receivers. If models are computed with very resistive layers or models with

highly resistive basements, this can even apply to deep marine scenarios.

Fig. 8 shows this effect. It is the same model as in the previous section; however, for the adaptive gridding in the horizontal directions, $\rho_{\text{ave}} = 1 \Omega \text{ m}$ was used instead of $\rho_{\text{ave}} = 10\,000 \Omega \text{ m}$. Having the boundaries too near in the horizontal directions leads to worse results for most frequencies and entirely wrong results for high frequencies. Comparison with the 1-D result in the time domain shows that it is the airwave whose amplitude is heavily overestimated. It can be difficult to spot these errors in the time-domain result, as the response looks plausible. Only a comparison with the 1-D result reveals that it is actually wrong. A possibility to detect such problems for complicated cases, where there is no semi-analytical result to compare with, is to carry out a mesh-convergence test by computing two or more models, moving the boundary. When the responses stop to change, one can assume that the boundary is far enough. Another possibility is to look at the amplitudes close to the boundaries and ensure that they are small enough.

4.4 Induced polarization

The proposed Fourier-transform approach works independently of spatial complexity, as diffusive electromagnetic fields are smooth functions of both frequency and time. An interesting test it to see whether this applies as well for dispersive media where the model parameters depend on frequency, as is the case in induced polarization. To test this we use the Cole–Cole model (CCM, Cole & Cole 1941) which, written in terms of conductivities instead of electric permittivities as by the original authors, is given by

$$\sigma(\omega) = \sigma_{\infty} + \frac{\sigma_0 - \sigma_{\infty}}{1 + (i\omega\tau)^c}, \quad (5)$$

Figure 8. Same model as used for Fig. 6, but with the horizontal boundaries not far enough. Although the resulting time-domain result looks plausible, the comparison with the 1-D result shows that it significantly overestimates the amplitude of the airwave.

where σ_0 and σ_∞ refer to the low-frequency and high-frequency conductivity values, respectively, τ is the central relaxation time (s) and c the CCM exponent describing the broadness of the relaxation time distribution. Note that this model slightly differs from the one phrased in terms of resistivities given by Pelton *et al.* (1978), see, for example Tarasov & Titov (2013).

Commer *et al.* (2017) have shown that time-domain IP modelling can be difficult, particularly for low c -values. Modelling frequency-dependent models in the frequency domain is, for any given frequency, not different from modelling a non frequency-dependent model. However, the interesting point is to see how the Fourier transform behaves to obtain time-domain data from frequency-domain responses. As an example, we modelled the response for a land case with air above a half-space of 1 S/m. Within the half-space is a dispersive, 100-m-thick layer at a depth of 300 m, with $\sigma_0 = 1.0$ S/m and $\sigma_\infty = 1.25$ S/m; $\tau = 1$, and we vary c from 0 (no IP) to 1.0, 0.75, 0.5 and 0.25. The x -directed step-off source and the receiver at an inline offset of 900 m are placed on the surface. The real and imaginary frequency-domain responses for this model are shown in the left-hand column of Fig. 9. It can be seen that the effect of the IP layer in the frequency domain is that the real part starts to deviate from the DC value at lower frequencies. The imaginary part decays slower towards zero when decreasing the frequency, and it deviates from a power-law function at lower frequencies as well. This indicates that we need lower frequency content for the Fourier transform, the smaller c becomes.

For the Fourier transform, we used the 601-point sine-cosine filter from Key (2009). This is a very long filter, but it seemed to be by far the best filter for low c -values. However, designing a new filter particularly for IP problems might be a better approach. The selected frequency thresholds were chosen at $f_{\min} = 10^{-4}$ Hz and $f_{\max} = 500$ Hz, and within the thresholds we only computed every 6th frequency. Using these values only 27 frequencies have to be

computed, instead of the 747 frequencies required by the actual transform. To obtain the step-off response we computed the step-on response and subtracted it from the DC value at 1000 s (without the need for any additional frequencies).

4.5 3-D model

The last example consists of a resistive 3-D block embedded in the lower of two half-spaces, as depicted in Fig. 10. The target has resistivity $\rho_{tg} = 100 \Omega \text{ m}$, the upper half-space corresponds to seawater with $\rho_{sea} = 0.3 \Omega \text{ m}$, and the lower half-space is the background with $\rho_{bg} = 1 \Omega \text{ m}$. The source is a 100-m-long, x -directed dipole at the origin, 50 m above the seafloor, and we are using a step-off source function. The x -directed inline receiver is at an offset of 2 km. The dimension of the target cube is $1.1 \times 1.0 \times 0.4$ km, located 300 m below the seafloor in the centre between source and receiver.

For the comparison, we use the open-source code SimPEG (Cockett *et al.* 2015), which is a framework for modelling and inversion of geophysical data such as gravity, magnetics and CSEM. It has Maxwell's equations implemented in both the frequency and time domain. As such we can compare our result computed in the frequency domain followed by a Fourier transform to a result computed directly in the time domain. A principal difference between SimPEG and emg3d is that the former has various direct solvers implemented, whereas the latter is an iterative multigrid solver. The 3-D model is therefore a rather small example in order to be able to run it on our test machine, as the memory requirement by the direct solver would otherwise be too high. There are not many options out there of open-source time-domain 3-D CSEM codes, SimPEG being the one we found to be suitable. A step-off response was chosen as this is the response currently implemented in it.

Figure 9. IP example for different c values, frequency-domain response on the left and time-domain response on the right. The frequencies indicated by the blue vertical bars in the lower-left figure were used for our Fourier transform approach, requiring only the computation of 27 frequencies instead of the full 747 frequencies required by the digital linear filter method.

Figure 10. 3-D block embedded in the lower of two half-spaces. The 100-m long, x -directed dipole source is located 50 m above the seafloor at the origin, and the receiver is on the seafloor at an offset of 2 km.

The model was discretized with $100 \times 100 \times 100$ m cells in the survey domain D_s . For the time-domain model, 14 cells in x -direction and 12 cells in y - and z -directions were used on both sides with a stretching of 1.3 for the total computation domain D_c , which yields a mesh of 58 344 cells. The time-steps start at 0.1 s and are: 21×0.01 s, 23×0.03 s, 21×0.1 s, 23×0.3 s, covering exactly the desired range of 0.1–10 s. For the frequency-domain model, the mesh is generated frequency-dependent as in the previous examples, with a maximum stretching of $\alpha_c = 1.5$. This results in meshes between 18 432 cells for the highest frequencies and 76 800 cells for the lowest frequencies. The required frequencies were obtained by using the FFTLog with five points per decade, which results in 20 frequencies between 0.001 and 8 Hz. The actual transform was carried out with the 201-point sine-cosine filter from Key (2009).

The results are shown in Fig. 11: In (a) the 1-D background responses and the relative error using the semi-analytical result and in (b) the responses including the target. The background comparison shows that both 3-D codes do an acceptable job with a relative error of a few percents at most; the result obtained through transformation seems to be better at early times. The reason is probably the implemented backward Euler scheme in the time-domain code that has an error of order one in time. We cannot compare the errors for the response that includes the target for lack of an analytical solution. The 1-D background model is only included to show that there is a significant response from the target. We therefore show the normalized difference (NRMSD) between the two responses R_1 and R_2 as a percentage, where $\text{NRMSD}(\%) = 200|R_1 - R_2|/(|R_1| + |R_2|)$. The NRMSD between the two codes is below 1 % everywhere except for early times. Both codes took roughly 4–5 min to compute the two models (single thread). However, in this particular comparison, the main difference in runtime is not frequency-domain computation versus time-domain computation, but iterative solver versus direct solver.

5 CONCLUSIONS

We have shown a method to minimize the required frequencies and their range for the computation of time-domain CSEM data with a frequency-domain code. This can significantly reduce the computation time and makes time-domain CSEM modelling with a frequency-domain code competitive given a robust frequency-domain solver, a frequency-dependent gridding function that minimizes the required cells, and a Fourier transform that works on a logarithmic scale. Fast layered modelling can be used to design the required frequency range, as the Fourier transform does

Figure 11. Responses for the model outlined in Fig. 10 using time-domain and frequency-domain computations, and for the layered background also the semi-analytical result. In the lower plot of (a) the relative error (%) is shown in comparison to the semi-analytical result and in (b) the normalized difference (%) between the two 3-D codes.

not know about the dimensionality of the underlying model. 20 frequencies or less are usually sufficient for a wide range of offsets. The values for lower frequencies can be interpolated using PCHIP knowing that the imaginary part goes to zero for zero frequency. The values for higher frequencies can be set to zero, as we can neglect their influence. And values for frequencies in-between the computed ones are best obtained with a spline interpolation. The actual transform can be carried out with either the DLF method or FFTLog, where the latter one requires usually much fewer frequencies to be interpolated. We have demonstrated the idea of our Fourier transform method on CSEM data transformed from the frequency domain to the time domain. However, it could equally be applied to the transform from the time domain to the frequency domain and to other methods with similar characteristics. We believe that our proposed improvements to the previously published methods makes simulating results in one domain obtained through computations in the other domain followed by a transformation a viable alternative. The methodology is relatively simple to implement and has therefore the potential to expand the capability of any existing code to an additional dimension.

ACKNOWLEDGEMENTS

We would like to thank the editor Ute Weckmann and assistant editor Louise Alexander as well as the reviewers Michael Commer and Chester Weiss for many helpful comments, which improved this manuscript considerably. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 728053-MarTERA. An early idea of this manuscript was presented by Werthmüller & Slob (2020).

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available at Zenodo at doi: 10.5281/zenodo.4710820. The data includes the scripts and instructions to reproduce all results and figures.

REFERENCES

- Andréis, D. & MacGregor, L., 2007. Time-domain versus frequency-domain CSEM in shallow water, in *SEG Technical Program Expanded Abstracts*, pp. 643–647, Society of Exploration Geophy, 10.1190/1.2792500.
- Avdeeva, A., Commer, M. & Newman, G.A., 2007. Hydrocarbon reservoir detectability study for marine CSEM methods: time domain versus frequency domain, in *SEG Technical Program Expanded Abstracts*, pp. 628–632, Society of Exploration Geophysicists, 10.1190/1.2792497.
- Cockett, R., Kang, S., Heagy, L.J., Pidlisecky, A. & Oldenburg, D.W., 2015. SimPEG: an open source framework for simulation and gradient based parameter estimation in geophysical applications, *Comput. Geosci.*, **85**, 142–154, 10.1016/j.cageo.2015.09.015.
- Cole, K.S. & Cole, R.H., 1941. Dispersion and absorption in dielectrics I. Alternating current characteristics, *J. Chem. Phys.*, **9**, 341–351, 10.1063/1.1750906.
- Commer, M. & Newman, G., 2004. A parallel finite-difference approach for 3D transient electromagnetic modeling with galvanic sources, *Geophysics*, **69**(5), 1192–1202, 10.1190/1.1801936.
- Commer, M. & Newman, G.A., 2006. An accelerated time domain finite difference simulation scheme for three-dimensional transient electromagnetic modeling using geometric multigrid concepts, *Radio Sci.*, **41**(3), 10.1029/2005RS003413.
- Commer, M., Petrov, P.V. & Newman, G.A., 2017. FDTD modelling of induced polarization phenomena in transient electromagnetics, *Geophys. J. Int.*, **209**(1), 387–405, 10.1093/gji/ggx023.
- Connell, D. & Key, K., 2013. A numerical comparison of time and frequency-domain marine electromagnetic methods for hydrocarbon exploration in shallow water, *Geophys. Prospect.*, **61**(1), 187–199, 10.1111/j.1365-2478.2012.01037.x.

- Constable, S., 2010. Ten years of marine CSEM for hydrocarbon exploration, *Geophysics*, **75**(5), 75A67–75A81, 10.1190/1.3483451.
- Davydycheva, S., Druskin, V. & Habashy, T., 2003. An efficient finite-difference scheme for electromagnetic logging in 3D anisotropic inhomogeneous media, *Geophysics*, **68**(5), 1525–1536, 10.1190/1.1620626.
- Druskin, V. & Knizhnerman, L., 1994. Spectral approach to solving three-dimensional Maxwell's diffusion equations in the time and frequency domains, *Radio Sci.*, **29**(4), 937–953, 10.1190/geo2011-0088.1.
- Fritsch, F. & Carlson, R., 1980. Monotone piecewise cubic interpolation, *SIAM J. Numer. Anal.*, **17**(2), 238–246, 10.1137/0717021.
- Gehrmann, R.A.S., North, L.J., Graber, S., Sztikar, F., Petersen, S., Minshull, T.A. & Murton, B.J., 2019. Marine mineral exploration with controlled source electromagnetics at the TAG hydrothermal field, 26°N Mid-Atlantic Ridge, *Geophys. Res. Lett.*, **46**(11), 5808–5816, 10.1029/2019GL082928.
- Ghosh, D.P., 1971. The application of linear filter theory to the direct interpretation of geoelectrical resistivity sounding measurements, *Geophys. Prospect.*, **19**(2), 192–217, 10.1111/j.1365-2478.1971.tb00593.x.
- Girard, J.-F., Coppo, N., Wawrzyniak, P., Bourgeois, B., Baltassat, J.-M. & Gadalia, A., 2015. Resistivity imaging for geothermal exploration, using controlled-source EM where magneto-telluric is not applicable: Model and field study, in *Proceedings World Geothermal Congress: IRENA Meeting on Island Energy Transitions*, Fort de France, France, pp. 1–8.
- Haber, E., Ascher, U.M. & Oldenburg, D.W., 2004. Inversion of 3D electromagnetic data in frequency and time domain using an inexact all-at-once approach, *Geophysics*, **69**(5), 1216–1228, 10.1190/1.1801938.
- Hamilton, A.J.S., 2000. Uncorrelated modes of the non-linear power spectrum, *Mon. Not. R. Astron. Soc.*, **312**(2), 257–284, 10.1046/j.1365-8711.2000.03071.x.
- Heagy, L.J. & Oldenburg, D.W., 2019. Modeling electromagnetics on cylindrical meshes with applications to steel-cased wells, *Comput. Geosci.*, **125**, 115–130, 10.1016/j.cageo.2018.11.010.
- Hunziker, J., Thorbecke, J. & Slob, E., 2015. The electromagnetic response in a layered vertical transverse isotropic medium: a new look at an old problem, *Geophysics*, **80**(1), F1–F18, 10.1190/geo2013-0411.1.
- Jönsthövel, T.B., Oosterlee, C.W. & Mulder, W.A., 2006. Improving multigrid for 3-D electro-magnetic diffusion on stretched grids, in *ECCOMAS CFD 2006: Proceedings of the European Conference on Computational Fluid Dynamics*, Delft University of Technology, UUID: df65da5c-e43f-47ab-b80d-2f8ee7f35464.
- Johansen, S.E., Panzner, M., Mittet, R., Amundsen, H.E.F., Lim, A., Vik, E., Landrø, M. & Arntsen, B., 2019. Deep electrical imaging of the ultraslow-spreading Mohs Ridge, *Nature*, **567**(7748), 379–383, 10.1038/s41586-019-1010-0.
- Kang, S., Oldenburg, D.W. & Heagy, L.J., 2020. Detecting induced polarisation effects in time-domain data: a modelling study using stretched exponentials, *Explor. Geophys.*, **51**(1), 122–133, 10.1080/08123985.2019.1690393.
- Key, K., 2009. 1D inversion of multicomponent, multifrequency marine CSEM data: methodology and synthetic studies for resolving thin resistive layers, *Geophysics*, **74**(2), F9–F20, 10.1190/1.3058434.
- Lam, S.K., Pitrou, A. & Seibert, S., 2015. Numba: a LLVM-based Python JIT compiler, in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pp. 7:1–7:6, ACM, 10.1145/2833157.2833162.
- Maaß, F.A., 2007. Fast finite-difference time-domain modeling for marine-subsurface electromagnetic problems, *Geophysics*, **72**(2), A19–A23, 10.1190/1.2434781.
- Mulder, W.A., 2006. A multigrid solver for 3D electromagnetic diffusion, *Geophys. Prospect.*, **54**(5), 633–649, 10.1111/j.1365-2478.2006.00558.x.
- Mulder, W.A., 2020. *Numerical Methods, Multigrid*, pp. 1–6, Springer International Publishing, Cham, 10.1007/978-3-030-10475-7_153-1.
- Mulder, W.A., Wirianto, M. & Slob, E., 2008. Time-domain modeling of electromagnetic diffusion with a frequency-domain code, *Geophysics*, **73**(1), F1–F8, 10.1190/1.2799093.
- Pedersen, L.B., Bastani, M. & Dinesius, L., 2005. Groundwater exploration using combined controlled-source and radiomagnetotelluric techniques, *Geophysics*, **70**(1), G8–G15, 10.1190/1.1852774.
- Pelton, W.H., Ward, S.H., Hallof, P.G., Sill, W.R. & Nelson, P.H., 1978. Mineral discrimination and removal of inductive coupling with multifrequency IP, *Geophysics*, **43**(3), 588–609, 10.1190/1.1440839.
- Plessix, R.-E., Darnet, M. & Mulder, W.A., 2007. An approach for 3D multiresource, multifrequency CSEM modeling, *Geophysics*, **72**(5), SM177–SM184, 10.1190/1.2744234.
- Rochlitz, R., Seidel, M. & Börner, R.-U., 2021. Evaluation of three approaches for simulating time-domain electromagnetic data using the open-source software custEM, *Geophys. J. Int.*, submitted.
- Tarasov, A. & Titov, K., 2013. On the use of the Cole–Cole equations in spectral induced polarization, *Geophys. J. Int.*, **195**(1), 352–356, 10.1093/gji/ggt251.
- Um, E.S., Harris, J.M. & Alumbaugh, D.J., 2010. 3D time-domain simulation of electromagnetic diffusion phenomena: a finite-element electric-field approach, *Geophysics*, **75**(4), F115–F126, 10.1190/1.3473694.
- Ward, S.H. & Hohmann, G.W., 1988. Electromagnetic theory for geophysical applications, in *Electromagnetic Methods in Applied Geophysics*, Chapter 4, Vol. 1: Theory, Investigations in Geophysics no. 3, pp. 130–311, ed. Nabighian, M.N., SEG, 10.1190/1.9781560802631.ch4.
- Weiland, T., 1977. Eine Methode zur Lösung der Maxwellischen Gleichungen für sechskomponentige Felder auf diskreter Basis, *Archiv für Elektronik und Übertragungstechnik*, **31**(3), 116–120.
- Werthmüller, D., 2017. An open-source full 3D electromagnetic modeler for 1D VTI media in Python: empymod, *Geophysics*, **82**(6), WB9–WB19, 10.1190/geo2016-0626.1.
- Werthmüller, D. & Slob, E.C., 2020. Time-domain CSEM modelling using frequency- and Laplace-domain computations, in *Proceedings of the 82nd EAGE Annual Conference & Exhibition*, Jul 2020, Vol. 2020, pp. 1–5, 10.3997/2214-4609.202010319.
- Werthmüller, D., Ziolkowski, A. & Wright, D., 2014. Predicting controlled-source electromagnetic responses from seismic velocities, *Interpretation*, **2**(3), SH115–SH131, 10.1190/INT-2013-0153.1.
- Werthmüller, D., Key, K. & Slob, E., 2019a. A tool for designing digital filters for the Hankel and Fourier transforms in potential, diffusive, and wavefield modeling, *Geophysics*, **84**(2), F47–F56, 10.1190/geo2018-0069.1.
- Werthmüller, D., Mulder, W.A. & Slob, E.C., 2019b. emg3d: a multigrid solver for 3D electromagnetic diffusion, *J. Open Sour. Softw.*, **4**(39), 1463, 10.21105/joss.01463.
- Ziolkowski, A., Hobbs, B. & Wright, D., 2007. Multitransient electromagnetic demonstration survey in France, *Geophysics*, **72**(4), F197–F209, 10.1190/1.2735802.