

A Sign-Preserving Filter for Signal Decomposition

B. Cárdenas*, T.R. Davenne, S.D. Garvey

Dept. of Mechanical, Materials and Manufacturing Engineering, University of Nottingham, United Kingdom

Corresponding author's email: Bruno.Cardenas@nottingham.ac.uk

Abstract

There are optimization problems in which an improvement in performance or a reduction in cost can be attained if the input signal of the system is split into multiple components. Splitting the signal allows customizing the design of the system's hardware for a narrower range of frequencies, which in turn allows making a better use of its physical properties.

There exist applications that have very specific signal-splitting requirements, such as "counter-flow avoidance", that conventional signal processing tools cannot meet. Accordingly, a novel "Sign-Preserving" filter has been developed and is presented in this paper. The underlying algorithm of the filter is comprehensively explained with the aim of facilitating its reproduction and the aspects of its operation are thoroughly discussed. The filter has two key features (1) it separates a discrete signal \mathbf{a} into two components, a mostly low-frequency signal \mathbf{b} and a predominantly high-frequency signal \mathbf{c} such that the sum of \mathbf{b} and \mathbf{c} replicates exactly the original signal \mathbf{a} and more importantly (2) the signs of the two output signals are equal to the sign of \mathbf{a} at all times.

The paper presents two case studies which demonstrate the use of the Sign-Preserving filter for the optimization of real life applications, in which counter-flow must be avoided: the hybridization of the battery pack of an electric vehicle and the parallelization of a packed bed thermal energy store.

Keywords

Signal splitting; load optimization; discrete signal processing; frequency based optimization; counter-flow avoidance;

List of Contents

1. Introduction
2. Motivation
3. The signal filtering process
4. Aspects of the operation of the filter
 - 4.1 Filtering a noisy signal
 - 4.2 Filtering a square wave
 - 4.3 Discussion on phase and signal energy
5. Examples of applications
 - 5.1 Optimization of an electric vehicle's battery pack
 - 5.2 Optimization of a packed bed thermal energy store
- 6 Concluding remarks
- 7 References

image and video processing and enhancement [3,4], control systems, medical equipment [5,6], geophysics [7], oceanography and astronomy [8].

Numerous applications in the aforementioned areas make use of signal processing tools to split a driving (input) signal into two or more signals based on their frequencies because it would be highly impractical, or impossible in some cases, to fabricate physical components capable of handling the whole spectrum of frequencies contained in the original signal. By separating the input signal into multiple components the overall system's performance can be improved or even optimized because the capacities of the different devices in the system are better exploited.

There are a number of well-known, fully characterized and extensively tested filter designs such as the Butterworth [9], Chebyshev [10], Bessel [11], Linkwitz-Riley [12] and Savitzky-Golay filters [13]—among others—by means of which the splitting of a signal into two or more frequency-based groups can be done. There exists an ample literature addressing their characteristics (response, delay, gain, etc.), as well as their advantages, shortcomings and applications for which each of them is particularly appropriate.

However there are certain applications, such as the hybridization of energy storage systems for cost-reduction or performance improvement, that have very specific signal-splitting requirements which the existing signal processing tools cannot meet.

Currently there is no signal-decomposition technique available that allows splitting a signal into two frequency based components in such a way that the sign of the two resulting signals is the same than that of the input signal at any given point. Accordingly, this paper presents a novel time-domain

1. Introduction

The processing of a signal refers to the transformation of a signal containing a package of information into a different form of signal(s) —typically simpler—for its utilization [1]. This branch of knowledge is of paramount importance for numerous aspects of modern life. Given its importance a considerable amount of research is still devoted to the study and development of improved signal processing techniques.

There are many areas where the processing and filtering of signals is applied. The following can be highlighted: consumer electronics, telecommunications and computer networks [2],

Nomenclature

Symbol	Definition	Symbol	Definition
α	Aspect ratio (Packed bed example)	E_s	Signal energy
γ	Vector containing the basic shape of the wavelets	i	Interval within the signal length
$\Delta\phi$	Phase difference with respect to original signal	k	No. of harmonic
η	Factor for impeding numerical errors	L	Signal length (amount of points)
θ	Vector of angles between 0 and 2π	\mathbf{M}	Matrix containing several \mathbf{c} signals for singular value decomposition
λ	Vector containing a shifted (downwards) version of γ	n_{runs}	Number of runs of the filter
μ	Zero-mean normalization of wavelet λ	norm	Magnitude of a vector
\mathbf{v}	Version of μ wavelet with starting and ending points at zero.	p_0	Initial estimation of correction coefficient
ϕ	Phase	p_{max}	Value to calculate max. permissible correction coefficient
χ	Battery capacity ratio (EV example)	p_{1b}	maximum permissible correction coefficient for snippet of \mathbf{b}
ω	Frequency of work cycle (for packed bed example)	p_{1c}	maximum permissible correction coefficient for snippet of \mathbf{c}
\mathbf{a}	Original signal to be filtered	p_1	Overall maximum permissible correction coefficient
\mathbf{b}	Low frequency output signal	Q	Energy (heat) content of work cycle (Packed bed example)
\mathbf{b}_i	Snippet of signal \mathbf{b} in the interval i	\mathbf{S}	Diagonal matrix from singular value decomposition
\mathbf{b}_{i0}	Initial mod. of snippet of \mathbf{b}	sgn	Sign function
\mathbf{b}_{i1}	Temporary vector containing points whose sign changed upon correction	t	Time index or number of data point
\mathbf{b}_{i2}	Modified snippet of \mathbf{b} at interval i	W	Wavelet width in number of points
\mathbf{c}	High frequency output signal	W_{max}	Maximum wavelet width
\mathbf{c}_i	Snippet of signal \mathbf{c} in the interval i	x	Coefficient for the Fourier series
\mathbf{c}_{i0}	Initial mod. of snippet of \mathbf{c}	y	Coefficient for the Fourier series
\mathbf{c}_{i1}	Temporary vector containing points whose sign changed upon correction	Z	Ratio of signal energies
\mathbf{c}_{i2}	Modified snippet of \mathbf{c} at interval i		

based filter that satisfies the very particular condition of sign preservation throughout the signal length. The paper is intended to provide researchers and technologists a valuable tool for the optimization of systems of diverse kinds.

The paper is structured in the following way: section 2 presents the motivation behind the development of the Sign-Preserving filter. Section 3 discusses the mechanics of the filter. Here, the algorithm followed by the filter is comprehensively explained with the aim of facilitating its reproduction and use for solving different problems that have the same ‘‘sign-preservation’’ requirement. Section 4 provides further detail of some aspects of the operation of the filter such as signal energy and conservation of phase. Furthermore, this section discusses the applicability of the filter for different kinds of input signals such as noisy waves and pulses. Section 5 presents two case-studies that demonstrate the use of the Sign-Preserving filter for the optimization of real-life applications: the hybridization of an electric vehicle’s battery pack and the parallelization of a packed bed thermal energy store. Lastly, section 6 provides a summary of the key takeaways of the paper.

2 Motivation

The motivation for the Sign-Preserving filter will be explained by means of Figure 1. In the figure, the input (or driving) signal \mathbf{a} for a certain system is shown. The signal, as it can be seen,

comprises a mix of low and high frequencies; therefore the physical component of the system that will be driven by it must be capable of handling a wide range of frequencies.

An improvement in the overall performance of the system could be achieved if the input signal \mathbf{a} was split into two signals \mathbf{b} and \mathbf{c} in such a way that \mathbf{b} contained the low frequencies and \mathbf{c} comprised the remaining high frequencies. Both signals \mathbf{b} and \mathbf{c} will be inputs to the system and will drive two separate physical devices. The improvement in performance (or cost reduction in some cases) is owed to the fact that the components driven by the signals can be customized for the specific (and narrower) range of frequencies they will experience as opposed to having a device whose performance is compromised because it was designed for being capable of handling the whole spectrum contained in \mathbf{a} .

As mentioned in section 1, there are many different tools that can be used to separate a signal into two or more frequency components. Figure 1a shows an example of a decomposition of \mathbf{a} carried out through a Fourier analysis, where the low frequency components were grouped in \mathbf{b} whilst the high-frequency components were grouped in \mathbf{c} . It should be noted that this kind of splitting allows a perfect reproduction of the original signal (i.e. $\mathbf{b} + \mathbf{c} = \mathbf{a}$).

It can be seen in Figure 1a that the two resultant signals have different signs from each other in many points in time throughout the signal's length (e.g. $t=12$). This phenomenon is called "counter-flow". In many applications it is not a problem at all; however in some others, such as the examples presented in section 5, having counter-flow is highly undesirable.

Suppose for the sake of argument, that the signal a is the duty of an electric motor. During the periods when a is positive the motor is consuming electricity and producing mechanical work while during the periods when a is negative the motor acts as a generator, i.e. consumes mechanical work and produces electricity.

When the signal is split into two components b and c , as shown in Figure 1a, two smaller motors are used, one for each new signal. These motors will work in the same way described, as a motor during the positive periods and as a generator during negative periods. The two motors are expected to work in parallel and their combined output should be the same as the output of the single motor driven by a . This condition is met in the decomposition shown in Figure 1a.

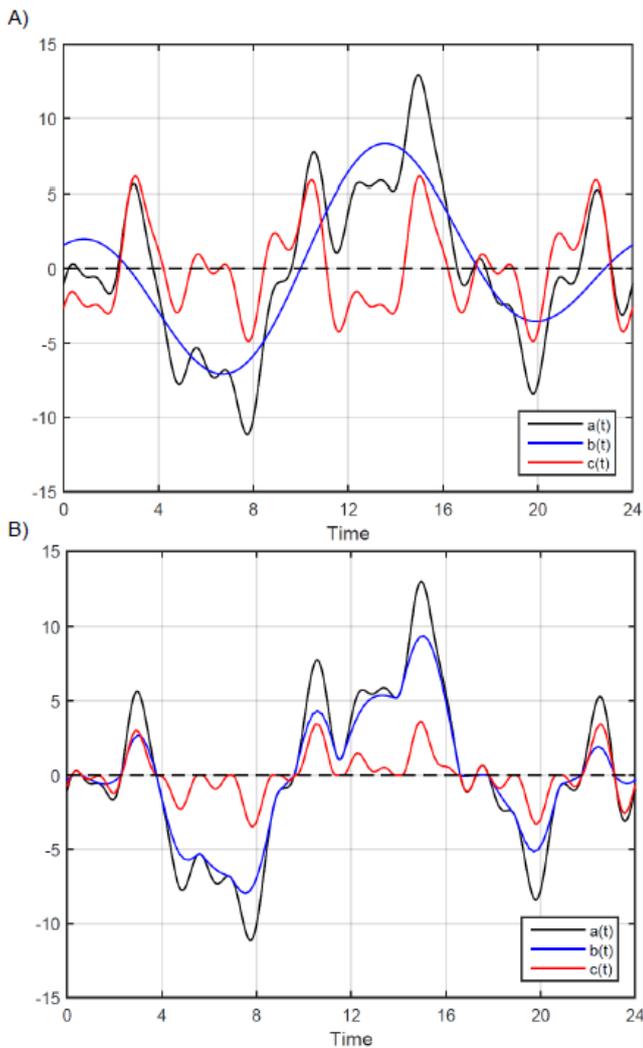


Figure 1. Examples of signal decompositions using a) Fourier analysis and b) Sign-Preserving filter

Nevertheless, there are many instances (e.g. $t=4$ and $t=12$) where counter-flow is found. In these cases, one motor is consuming electricity (producing mechanical work) while the other one is operating in reverse producing electricity. In $t=12$,

for example, the motor driven by b is consuming more electricity than what the motor driven by a would be consuming; this excess power is provided by the motor driven by c . Having counter-flow in the system means that the effects of b and c counter-act each other (instead of adding to each other) to produce the desired final effect, which is to replicate a .

The foregoing is just a very simple example of how in certain systems a split of the signals in a conventional way does not make much sense. Any of the existing signal-splitting techniques, either frequency-domain based tools such as the Butterworth, Chebyshev, Bessel and Linkwitz-Riley filters or time-domain based methods such as the Savitzky-Golay filter (or a more basic moving average approach) will produce output signals (b and c) with some degree of counter-flow.

Consequently, this paper proposes a filter by means of which a discrete time signal a can be split into two different components without incurring in the aforementioned counter-flow problem. The key-feature of this "Sign-Preserving Filter" is that at any given time the following two conditions are met:

$$a = b + c \quad (1)$$

$$\text{sgn}(a) = \text{sgn}(b) = \text{sgn}(c) \quad (2)$$

The component b contains mainly low-frequency (smooth) content while c is predominantly composed of high frequency (not-so-smooth) content. The filter, as Eqs. (1) and (2) indicate is designed to produce 2 outputs b and c ; however it can be reapplied on the resultant signal b to produce 3 outputs: b_1 , b_2 and C .

Figure 1b shows an example of a decomposition of the same signal a carried out through the Sign-Preserving filter. The big difference the Sign-Preserving filter has with other Fourier-based filters is that the condition described by Eq. (2) is true at all times.

The Sign-Preserving filter can be used in numerous optimization problems where an improvement in the performance of the system or a cost reduction can be achieved by splitting the driving signal but counter-flow must be avoided. The algorithm followed by the filter is explained in full detail in the subsequent sections.

3. The signal filtering process

This section discusses in depth the mechanics of the operation of the filter. The algorithm followed by the Sign-Preserving filter is shown in Figure 2.

The operation of the filter is based on passing a number of wavelets of different sizes throughout the length of the signal to be filtered. As the wavelets are passed, some non-smoothness is subtracted from the signal and that subtracted content is accumulated elsewhere.

There are two parameters that control the operation of the filter. The first is the maximum width of the wavelets (W_{max}) and the second is the number of runs (n_{run}). A run of the filter consists in passing twice through the signal, all the odd wave-widths contained in the interval between $W=3$ and $W=W_{max}$.

A larger number of runs will produce a smoother signal \mathbf{b} at the expense of a longer computation time.

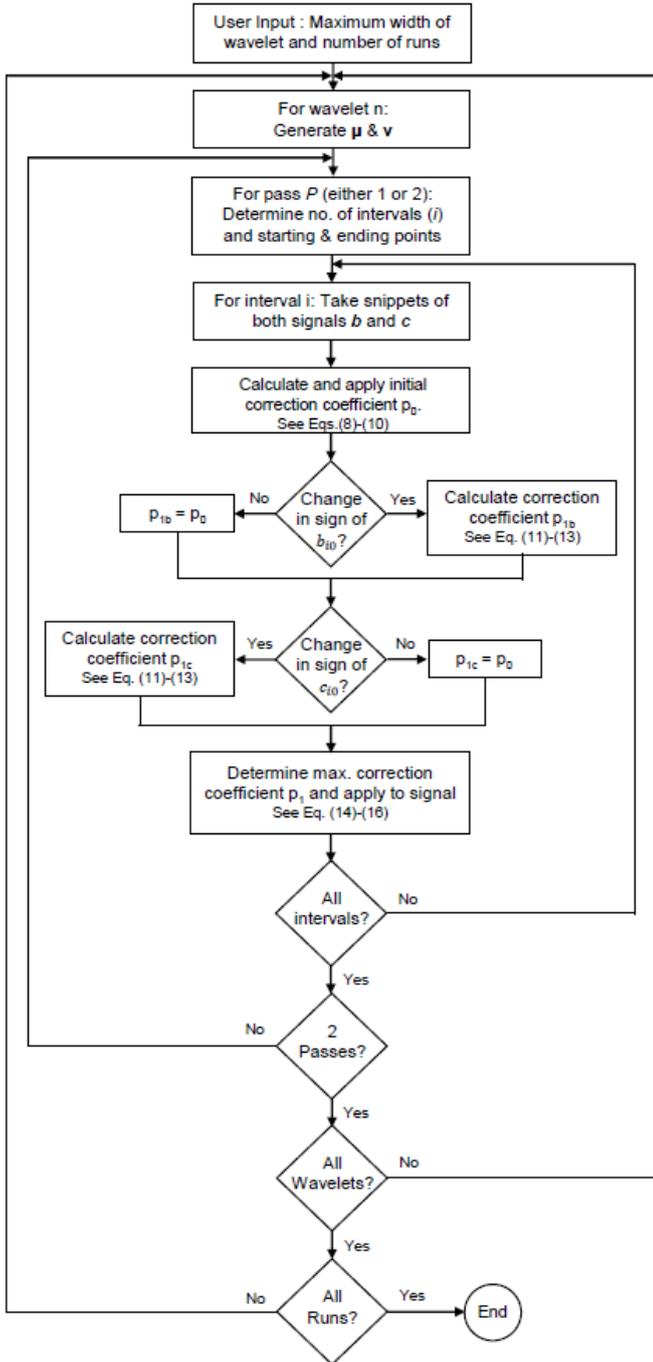


Figure 2. Algorithm followed by the Sign-Preserving filter

Every wavelet has a value of zero at both of its ends, as shown in Figure 3. The way the wavelets are generated and their shape will be explained in detail further ahead in the paper. The fact that the wavelets start and end having a value of zero means that they do not have any effect on those points. In order to modify those points, each wavelet is passed twice through the signal. During the second sweep the wavelet is offset by half of the wavelet-width. For example, if a wavelet with a $W=7$ started at point 1 of the signal during the first pass, it will start at point 4 during the second pass. The foregoing is graphically explained by Figure 4. In order for the wavelet's starting point during the second pass to be exactly at the middle point of the wavelet during the first pass, the length of the wavelet has to be an odd number of points. Therefore only odd values between $W=3$ and $W= W_{max}$ are considered.

The wavelets are generated in the following way: A vector of angles (θ) containing W equally spaced points between 0 and 2π is created. The basic shape of the wavelets (γ) is given by Eq. (3). The waves μ and ν , given by Eqs. (5) and (6) respectively, are normalized versions of the basic γ curve. These two waves are used by the filter during its operation to determine the amount of subtraction or “correction” to be applied to the signal at a given point.

$$\gamma = 1 - \cos(\theta) \quad (3)$$

$$\lambda = \gamma - \text{mean}(\gamma) \quad (4)$$

$$\mu = \frac{\lambda}{\sqrt{\lambda^T * \lambda}} \quad (5)$$

$$\nu = \frac{\gamma}{\sqrt{\mu^T * \gamma}} \quad (6)$$

Figure 3 shows examples of wavelets generated for a $W=3, 9$ and 51 . It can be seen that as aforementioned, the starting and ending points of the ν wavelet have always a value of zero, which is an important aspect of the filter's operation. It can also be seen that the wavelets (which define an interval of interest during the filter's operation) resemble a Hann Window; although within the Sign-Preserving filter they are used in a different way.

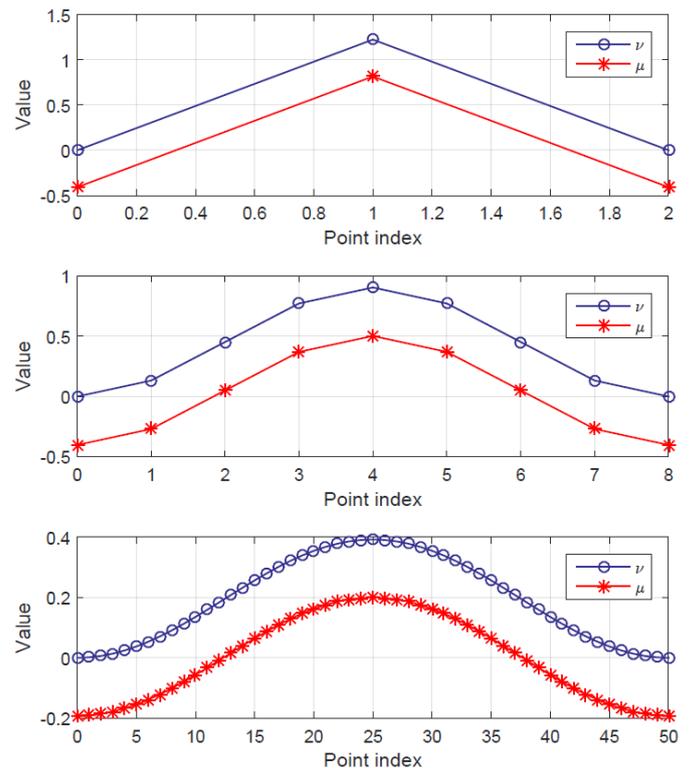


Figure 3. Wavelets μ and ν generated for a $W=3$ (top), $W=9$ (middle) and $W=51$ (bottom)

The two components of the signal \mathbf{b} and \mathbf{c} are initialized. These components have the same length (L) as \mathbf{a} . The signal \mathbf{c} which will hold the high frequency part is set to zero while \mathbf{b} is initialized as an identical copy of \mathbf{a} . As the filter progresses some of the content of \mathbf{b} will be removed (i.e. filtered) and stored in \mathbf{c} .

After the initialization of the two components the operation of the filter begins. The first pass is done with a 3-point wavelet, which is the smallest possible wavelet. The signal \mathbf{b} is divided into multiple intervals depending on the width of the wavelet in course. The number of intervals (i) is given by Eq. (7). For example, for a 3-point wavelet, the intervals go from points 1 to 3, from points 3 to 5, from points 5 to 7, and so on. Every interval is modified as the filter marches through the signal length.

$$i = \text{ceiling} \left[\frac{L}{W-1} \right] \quad (7)$$

In some cases the last interval is incomplete, in which case there is a wrap-around. The interval covers whatever points are left and continues at the beginning of the signal. For example, in a signal with 60480 points, the last interval for a 51-point wavelet would start at point 60451 and only has 30 points before reaching 60480 (the last point of the signal), so it extends to points 1-21.

The points of the signal that are used for the aforementioned “wrap-around” are not modified, as they have been already filtered at the start of the pass. The wrap-around is just used to complete the interval so that the filter can calculate the amount of “correction”, which is applied only to the remaining points at the end of the signal.

The above implies that the filter is not symmetric hence it cannot be applied in reverse (i.e. start at point n of signal and work towards point 1) because the correction calculated for the wrapped-around intervals will be different in both cases (forwards and backwards) due to the different shape of the signal at both of its ends. The outputs produced by the forward run will not be exactly equal to those generated by the reverse run but they will not differ greatly from each other. Nevertheless, there is no practical reason why the filter would need to be run backwards thus the asymmetry is not a problem.

It is worth highlighting that the intervals are independent from each other. The ending point of any given interval is the starting point for the next one. During a pass of a wavelet those points remain unaltered, because the starting and ending points of the wavelet are always zero. This independency offers the possibility to carry out operations in parallel, that is to say, within a pass of a wavelet all the intervals could be filtered simultaneously rather than sequentially, which offers the possibility to employ more computational power to reduce the calculation time.

As it was mentioned before, two passes are carried out for each wavelet. After the first sweep through the signal has been completed a second sweep is performed with an offset of half the wavelet width. For example, wavelets with a $W=9$ and $W=25$ start at point 1 during the first pass but during the second pass start at points 5 and 13 respectively, as Figure 4 shows. By doing so, with each wave-width a cosine and a sine wave are effectively being applied. The second pass (with an offset) is carried out to take care of the starting and ending points of the intervals that were not modified during the first pass.

Inside every interval, the algorithm only sees a snippet \mathbf{b}_i and \mathbf{c}_i of the complete signals \mathbf{b} and \mathbf{c} , which are determined by the width of the interval. For a given interval i , the algorithm

calculates, by means of Eq. (8), the correction coefficient (p_0) that would be applied to the signal snippets if there were no sign constraints in place. This coefficient is a mere initial estimate.

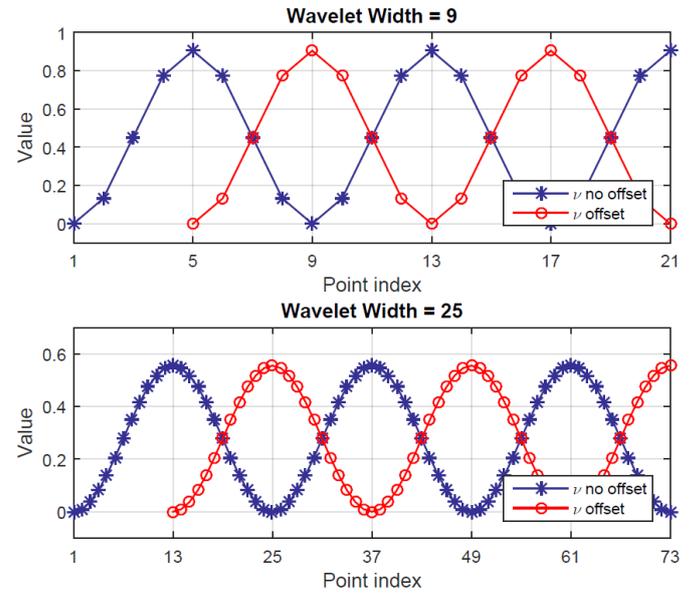


Figure 4. Comparison of the \mathbf{v} wavelet with a $W=9$ (top) and $W=25$ (bottom) during the first and second pass

$$p_0 = \boldsymbol{\mu}^T * \mathbf{b}_i \quad (8)$$

Both signals are modified based on the correction coefficient (p_0). The filter subtracts from the snippet of the low-frequency signal some non-smoothness and adds it to the snippet of the high-frequency signal, as stated in Eqs. (9) and (10).

$$\mathbf{b}_{i0} = \mathbf{b}_i - p_0 * \mathbf{v} \quad (9)$$

$$\mathbf{c}_{i0} = \mathbf{c}_i + p_0 * \mathbf{v} \quad (10)$$

Subsequently, a series of checks are carried out to determine if the correction applied caused a sign change within any one of the signal snippets. A sign change in either of the modified signal snippets means that the initial estimation for p_0 is too large and needs to be revised. The points of the low and high frequency signal snippets (\mathbf{b}_{i0} and \mathbf{c}_{i0} , respectively) whose sign changed (in either direction) are extracted and stored in smaller temporary vectors called \mathbf{b}_{i1} and \mathbf{c}_{i1} . Their corresponding entries of the \mathbf{v} wavelet are also extracted and stored in vectors called \mathbf{v}_{1b} and \mathbf{v}_{1c} .

The maximum permissible correction coefficients p_{1b} or p_{1c} that can be applied to the signals are obtained through Eqs. (11)-(13)

$$p_{\max b} = \min \left(\frac{|\mathbf{b}_{i1}|}{\mathbf{v}_{1b}} \right) \quad (11)$$

$$p_{\max c} = \min \left(\frac{|\mathbf{c}_{i1}|}{\mathbf{v}_{1c}} \right) \quad (12)$$

$$p_{1b,c} = \text{sgn}(p_0) * \min(|p_0|, p_{\max b,c}) \quad (13)$$

If no change of sign was observed in either of the modified signals (\mathbf{b}_{i0} or \mathbf{c}_{i0}) then the corresponding correction coefficient (p_{1b} or p_{1c}) is equal to the p_0 originally calculated by Eq.(8). The maximum correction coefficient that can be applied (p_1) is the minimum between p_{1b} and p_{1c} . This ensures that neither of the signals experiences a change in sign. It should be mentioned that either one of the signals can take a value of zero, zero is considered neutral and doesn't represent a change of sign. A factor η of 0.999 is introduced to avoid problems caused by numerical inaccuracy. If the full p_1 coefficient is applied, there is a risk that a sign change will still occur due to rounding errors. The two signals are modified as Eqs. (15) and (16) indicate.

$$p_1 = \min(p_{1b}, p_{1c}) \quad (14)$$

$$\mathbf{b}_{i2} = \mathbf{b}_i - p_1 * \eta * \mathbf{v} \quad (15)$$

$$\mathbf{c}_{i2} = \mathbf{c}_i + p_1 * \eta * \mathbf{v} \quad (16)$$

Once the new values (\mathbf{b}_{i2} and \mathbf{c}_{i2}) for the snippets of both signals have been calculated, these are installed in the full signals \mathbf{b} and \mathbf{c} in their respective positions. Following, the algorithm moves to the next interval where the filtering operations discussed above are repeated. After all intervals are completed for a given wavelet a second pass considering an offset in the starting point of the wavelets is carried out. When the two passes of a wavelet are finished the algorithm goes back to the early beginning and generates a new wavelet based on the next odd number of points (i.e. $W=W+2$). The loop repeats until all wavelets in the range between $W=3$ and W_{max} have been passed through the signal, at which point one run of the filter is completed. As aforementioned, several runs can be carried out (n_{runs}) to produce a smoother low-frequency component \mathbf{b} .

4 Aspects of the operation of the filter

In this section of the paper the operation of the Sign-Preserving filter is extensively reviewed. The applicability of the filter to different types of signals is demonstrated and some particularities of its behaviour—such as the conservation of phase and reduction of signal energy—are discussed and thoroughly analysed.

4.1 Filtering a noisy signal

Consider the signal \mathbf{a} described by the Fourier series given by Eq. (17), where $\boldsymbol{\theta}$ is a vector of 60480 points evenly spaced between 0 and 2π . The coefficients x and y for the series can be found in Table 1, for all the harmonics that do not appear on the table (e.g. $k=2, 7, 25$, etc.) the coefficients x and y have a value of zero.

$$\mathbf{a} = \sum_{k=1}^{113} x_k \cdot \cos(k\boldsymbol{\theta}) + y_k \cdot \sin(k\boldsymbol{\theta}) \quad (17)$$

Figure 5 shows an example of the operation of the filter. In this example the filter carried out 50 runs through the signal \mathbf{a} (described in Table 1) with a maximum wavelet width (W_{max}) of 79. As it can be seen, the filter has effectively created a

smoother signal (\mathbf{b}) while the high-frequency components removed have been stored in a separate signal (\mathbf{c}). It is also possible to observe that the two output signals produced have the same sign as \mathbf{a} at all times, which satisfies Eq. (2).

Table 1. Values for the α and β coefficients in Eq. (17)

k	x	y	Magnitude	Phase (φ)
1	-0.556	0.541	0.776	135.783°
4	-0.476	1.034	1.138	114.719°
9	-0.375	-0.539	0.657	-124.828°
11	-0.194	-0.141	0.240	-143.99°
37	-0.219	0.025	0.220	173.488°
79	-0.053	0.122	0.133	113.481°
113	0.089	0.047	0.101	27.838°

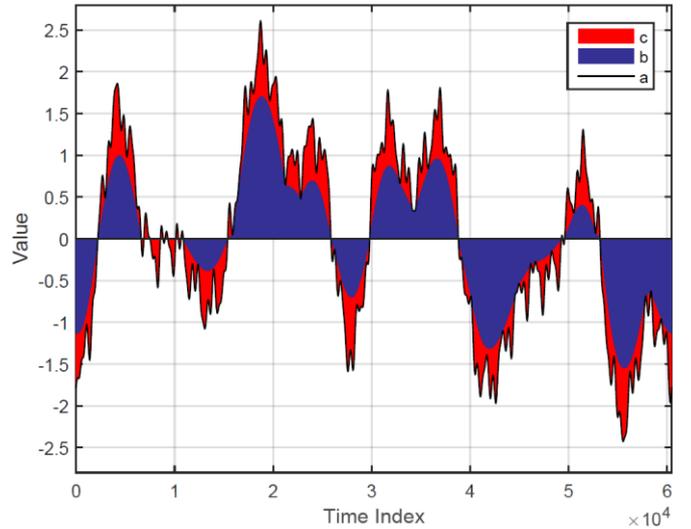


Figure 5. Example of the operation of the Sign-Preserving filter: Splitting of \mathbf{a} with a W_{max} of 79 and 50 runs

The Sign-Preserving filter has two control variables, the number of runs (n_{run}) and the maximum wavelet width to be used (W_{max}). Figure 6 shows a comparison of the results produced by the filter when 15, 30 and 100 runs are carried out with W_{max} values of 49, 125 and 199.

It is important to remember that W_{max} can only take odd values and its size is relative to the number of points (L) contained in the discrete signal (\mathbf{a}) being filtered.

As expected, a larger n_{run} yields a smoother low-frequency signal. However, the maximum wavelet width has a stronger effect on the performance of the filter. When a small value for W_{max} is used the filter can only remove a very small amount of “non-smoothness” on every run because it does fewer passes per run (proportional to the number of wavelets). Conversely, when a larger value for W_{max} is considered, more wavelets are used per run therefore the filter is able to subtract larger amounts of high-frequencies in every run and produce a much smoother curve in less iterations.

As aforementioned, to reach a certain result (i.e. level of smoothness) more runs will be required if a smaller number of wavelets are used per run. To exemplify this we'll use the signal \mathbf{a} described by Eq. (17). A way of measuring how much content has been taken from \mathbf{b} and put into \mathbf{c} is to measure the magnitude of the resultant signal, which is calculated by means of Eq. (18).

$$\text{norm}(\mathbf{b}) = \sqrt{b_1^2 + b_2^2 + \dots + b_{n-1}^2 + b_n^2} \quad (18)$$

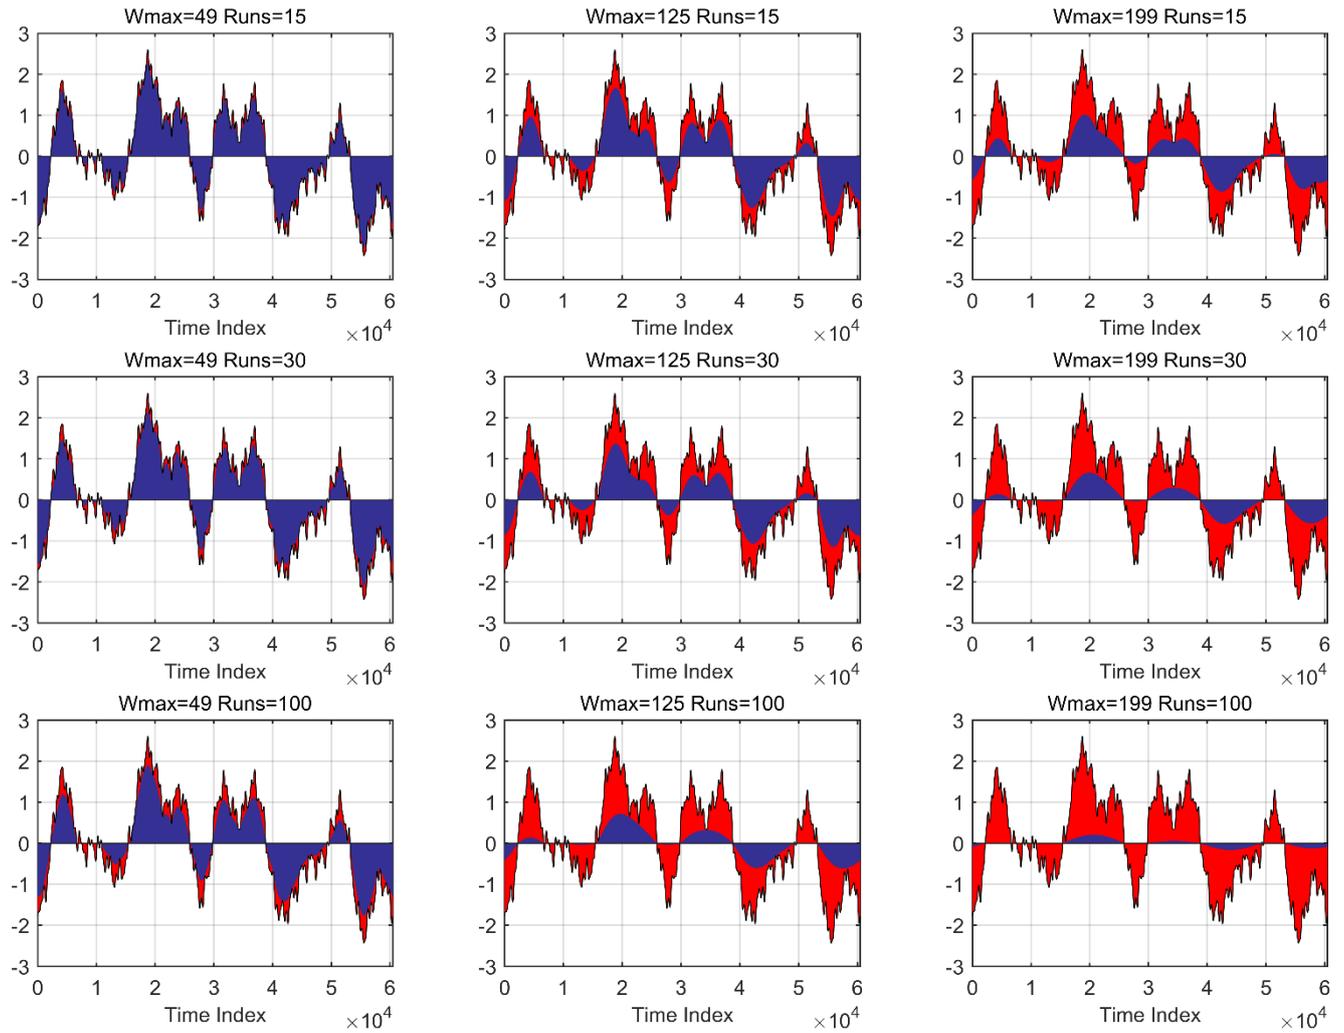


Figure 6. Signal split into \mathbf{b} (blue) and \mathbf{c} (red) using different combinations of control parameters: W_{max} and n_{runs}

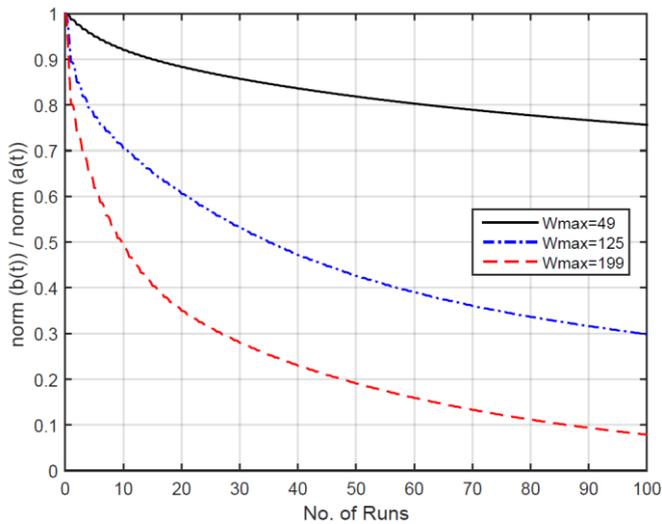


Figure 7. Reduction of magnitude of \mathbf{b} with respect to original signal \mathbf{a} as more runs of the filter are carried out.

For example, to produce a low-frequency signal \mathbf{b} whose magnitude is 50% of the magnitude of the original signal \mathbf{a} , a filtering operation considering a $W_{max}=199$ would require 10 runs. If a $W_{max}=125$ is used then the filtering operation will take 35 runs whereas if W_{max} is set to 49, the filter will need 569 runs to reach the specified value. Figure 7 shows the

reduction in magnitude of \mathbf{b} with respect to the original signal \mathbf{a} for W_{max} values of 49, 125 and 199.

It should be kept in mind that the values quoted are specific for the example used; however the behaviour depicted is a general trend. When larger values of W_{max} are used, the filter will shave more content from the signal with every run. The foregoing comparison does not suggest that using a large number of wavelets per run is necessarily good. The end goal of the filter will vary depending of the application for which it is employed and in some cases a precise or subtle filtering may be required, which cannot be done when W_{max} is set to a large value (relative to the number of data points in the signal). A good rule of thumb for defining W_{max} is that the maximum width of the wavelets should be $\leq 1\%$ of the total length of \mathbf{a} .

A question that arises is if it is possible to obtain the exact same filtered signal \mathbf{b} using a small n_{run} with a large W_{max} than with a large n_{run} and a small W_{max} . Except for the extreme case when the signal is completely flattened out (which needless to say is not at all the purpose of the filter) the answer is no, it is not possible. Several filter operations have been carried out (as shown in Figure 6) considering values for W_{max} of 49, 125 and 199 and values for n_{run} of 3, 15 and 30. With the results obtained a singular value decomposition in the form of $\mathbf{M}=\mathbf{USV}^*$ was carried out, where \mathbf{M} is a 60480 x 9

matrix containing the high-frequency signals (\mathbf{c}) produced by the filter. None of the singular values contained in the diagonal matrix \mathbf{S} are zero (or close to zero) which indicates a clear independency between the two control parameters of the filter. The above means that in the 2D space containing all the different possible combinations of W_{max} and n_{run} , there are no two combinations of values that yield the exact same result.

Another interesting feature (or quirk) of the Sign-Preserving filter that should be pointed out is the fact that, unlike filters based on a cut-off frequency such as the Butterworth and Chebyshev low/high pass filters, the Sign-Preserving filter does not have a stopping condition other than the specified number of runs. Hence, given enough time (i.e. a sufficiently large number of runs) the filter will smoothen a signal down to a straight line at zero. In other words, the filter will empty the vector containing the low-frequency signal \mathbf{b} and will eventually create a copy of \mathbf{a} in the vector containing the high-frequency signal \mathbf{c} . However, the key point is that the operation of the filter is a selective process: the highest-frequencies found in each pass are removed first and the low frequencies are left behind for a subsequent pass. The foregoing can be demonstrated via a Fourier analysis of the results produced by the filter.

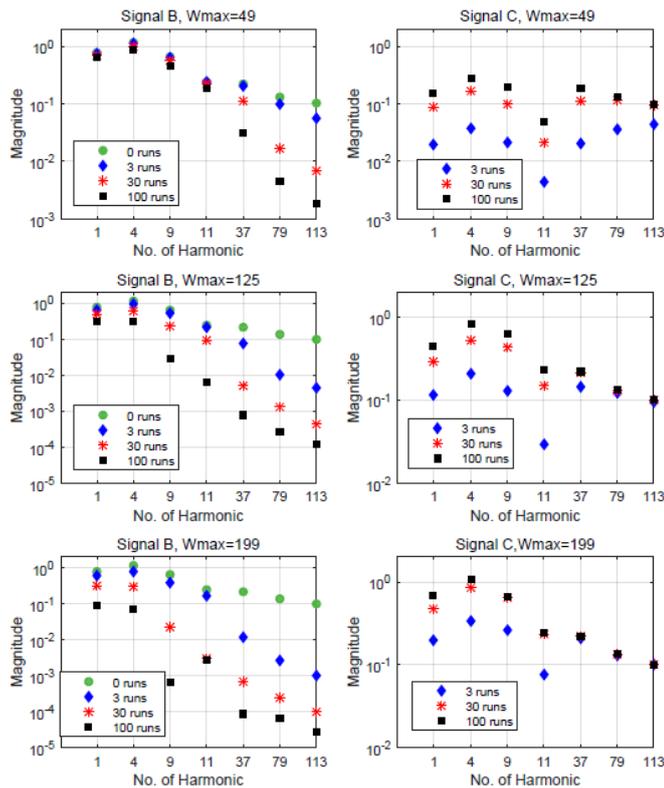


Figure 8. Magnitudes of the relevant frequency components of \mathbf{b} (left) and \mathbf{c} (right) as more runs of the filter are carried out.

Figure 8 shows, from a frequency-domain point of view, the result of the filter’s operation on the signal \mathbf{a} , described by Eq. (17). The filter was applied using three different values for W_{max} (49,125 and 199) and with each of them 3, 30 and 100 runs were carried out. Figure 8 shows the magnitudes of the frequency components of the output signals \mathbf{b} and \mathbf{c} . The time-domain results of these decompositions are shown in Figure 6. It can be easily seen in the plots for the \mathbf{b} signal (left) that the magnitude of the fast harmonics is considerably reduced

whilst the low-frequency components remain practically unaltered, which demonstrates the selectivity of the Sign-Preserving filter during its operation.

It can also be seen that when a large w_{max} is used together with a large number of runs (i.e. $w_{max} = 125$ and $n_{run} = 100$ or $w_{max} = 199$ and $n_{run} \geq 30$) the frequency spectrum of the \mathbf{c} signal, which should ideally comprise mostly high-frequency content, starts to resemble that of the \mathbf{a} signal. In other words, given enough time \mathbf{c} becomes a copy of \mathbf{a} .

Moreover, Figure 8 helps to further illustrate the fact that when a small value for W_{max} is used, the filter is only capable of removing a very small amount of non-smoothness from the signal \mathbf{b} in every run. This is not necessarily an inconvenience as in some situations the ability to perform a fine adjustment may be useful which could not be achieved by a run consisting of several wavelet widths (i.e. a large W_{max}).

An important question to address is how to achieve a good decomposition and how to choose good values for w_{max} and n_{run} . The filter’s aim is to split a signal \mathbf{a} into two signals, a predominantly low-frequency signal \mathbf{b} and a predominantly high-frequency signal \mathbf{c} , which must have the same sign at all times and their sum must replicate exactly the original signal \mathbf{a} . As it has been shown in the paper, the filter is successful in doing that. Without an application (an optimization problem) a given signal split cannot be judged; as long as it meets the criteria established by Eqs. (1) and (2) then it is a good decomposition. One combination of $w_{max} + n_{run}$ is no better than another.

When using the filter as a tool for solving a specific optimization problem (such as the ones presented in section 5) a two-dimensional search space has to be created with the two control parameters: w_{max} and n_{run} . Each combination of values for those two parameters will produce a different signal-split that is, a pair of signals \mathbf{b} and \mathbf{c} . Every pair of signals will then be used as the input to a model, which is specific of the particular problem to be solved; the model is a “black box” to the filter. The model could be for example, a pair of equations that describe the cost of two devices that will be driven by the two signal-components. In this case the best signal split (combination of $w_{max} + n_{run}$) is the one that yields the lowest total cost. In other words, the best values for w_{max} and n_{run} are given by the application for which the filter is being used and can only be determined after an exhaustive exploration of their combinational space.

4.2 Filtering a square wave

In this subsection the use of the Sign-Preserving filter for filtering impulses and square waves of different lengths (number of points) is demonstrated. Square waves are completely different from the noisy waves (see Figure 5) that would typically be encountered in the applications for which the Sign-Preserving filter has been developed; however they are useful for further demonstrating the capabilities of the Sign-Preserving filter as they have a rich frequency spectrum.

Figure 9 shows the split of three different square waves with lengths of 1, 1000 and 30,000 points. The waves have a constant value of -1 and rise to 1 at a time index of 10,000.

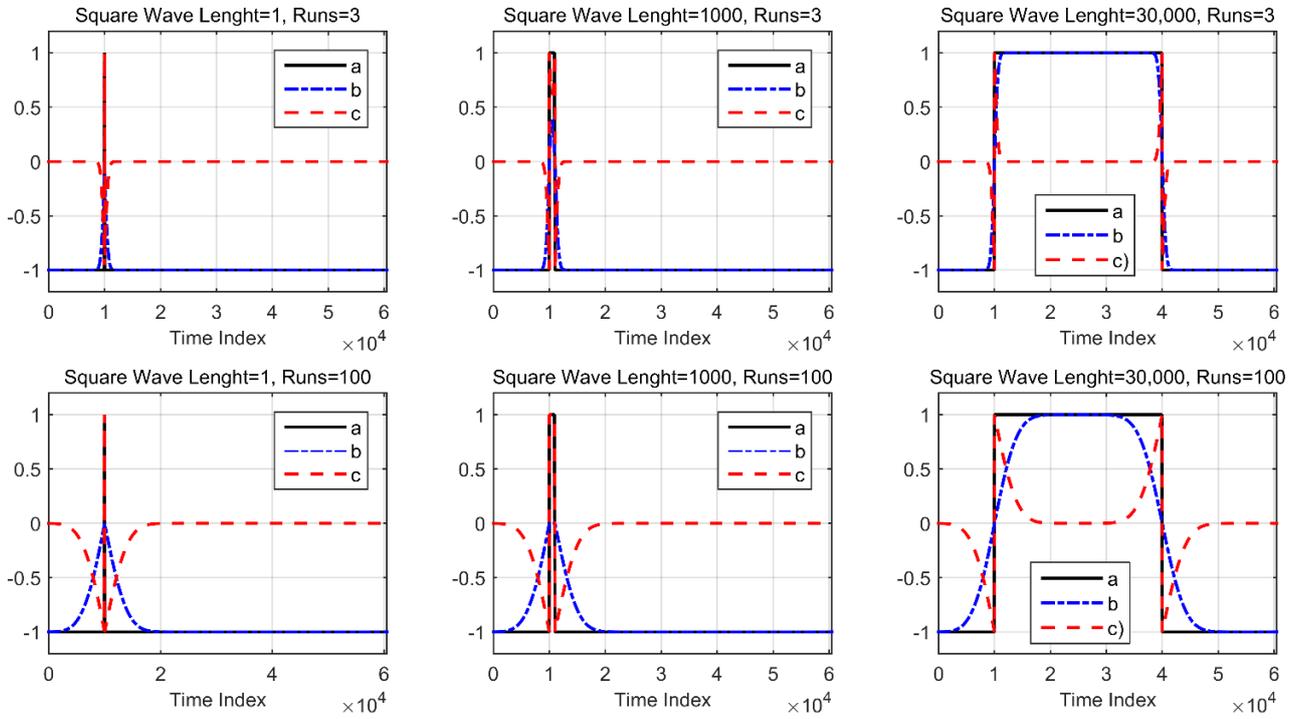


Figure 9. Split of square waves of different length into b and c with a fixed value for $W_{max}=125$

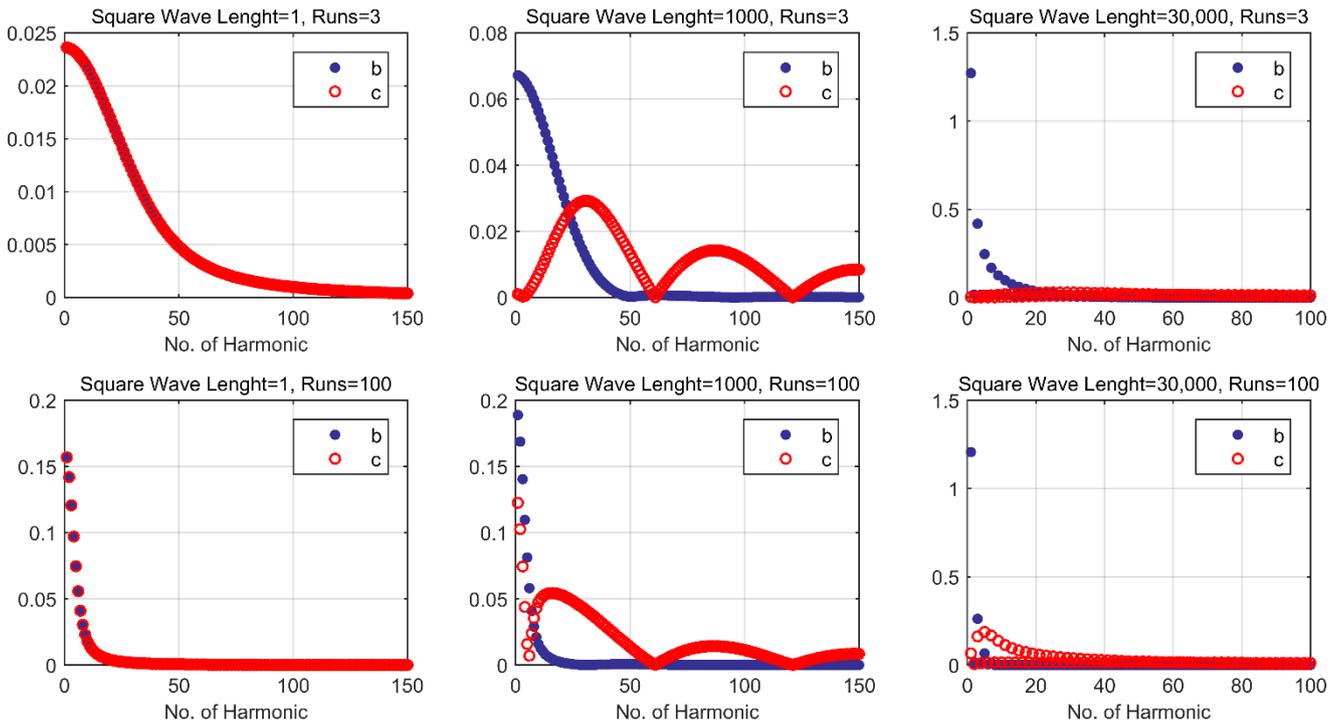


Figure 10. Frequency spectrum of signals b and c produced from the decomposition of square waves of different lengths

The peak value of 1 is maintained for 1, 1000 or 30,000 points, depending on the case, and after that it returns to -1. The filter was used on each of the three waves having a fixed value for $W_{max}=125$. Results produced after 3 and 100 runs are shown.

The Sign-Preserving filter is successful in splitting the square waves into a mostly low frequency signal b and a mostly high-frequency signal c that retains the sharp transitions. More importantly, the condition that the sign of the two components created should be the same as that of a at all times is met.

It can be seen—perhaps more evidently in the case of a 30,000 point-long square wave—that as more runs are carried out b

starts to widen at the base in order to smooth out the abrupt rise of the square wave from a negative to a positive value.

Figure 10 shows, complementary to the time-domain based results presented in Figure 9, the frequency spectrums of the low and high frequency signals generated from the split of the three different square waves. In the case of the 1-point wave (or impulse), it can be seen that as more runs of the filter are carried out, the magnitudes of the frequency components of both signals start to increase, which in a way is the opposite effect to what was observed in the examples presented in section 4.1. This occurs because as more passes of the filter are carried out, both signals start to cover a wider spread in the

time domain. One may also observe that the signal \mathbf{b} becomes increasingly a low-frequency signal as the filter completes more runs. The magnitude of the slow-harmonics increases considerably while some fast harmonics (e.g. 30-100) practically disappear from it. Another interesting phenomenon to point out for the case of the 1-point wave is that the magnitudes of the frequency spectrums of both signals, \mathbf{b} and \mathbf{c} , are very similar to one another.

A similar behaviour can be seen in the case of the 1000-point square wave. The magnitudes of the slow harmonics of \mathbf{b} increase significantly as more runs of the filter are carried out while the fast-harmonics become negligible, becoming progressively a predominantly low-frequency signal. Interestingly, the magnitudes of some slow-harmonics in \mathbf{c} increase noticeably as well; nevertheless \mathbf{c} remains being a (mostly) high-frequency signal.

4.3 Discussion on phase and signal energy

Most of the applications for which the Sign-Preserving filter could be used (such as the ones presented in section 5) are only concerned with avoiding counter-flow in the output signals. Notwithstanding, this section discusses the behaviour of the Sign-Preserving filter in terms of signal energy (E_s) and phase (φ), with the aim of providing a broader overview of the characteristics of the signal splitting tool that allows the filter to be used in different unforeseen applications that require for instance, a certain degree of control over the phase of the output signals, in addition to the sign preservation feature.

A case study was carried out with 10,000 randomly generated waves. The random waves are composed of the same harmonics (1, 4, 9, 11, 37, 79 and 113) as the \mathbf{a} signal described by Eq. (17). The values for the x and y coefficients that determine the magnitude and phase of each of them are randomly generated within certain specified ranges. The ranges have been defined to ensure that the waves created have strong low-frequency components with some added high-frequency noise, which resembles the signals that would most likely be encountered in the envisaged real-life applications of the Sign-Preserving filter. The ranges for the values that x and y may acquire are given by Table 2. Figure 11 shows some examples of the randomly generated noisy-signals \mathbf{a} used for the phase and signal energy study.

Table 2. Ranges of values for the x and y coefficients of the random waves generated

k	x	y
1,4	$\{-5 \leq x \leq -3.5\} \cup \{3.5 \leq x \leq 5\}$	$\{-5 \leq y \leq -3.5\} \cup \{3.5 \leq y \leq 5\}$
9,11	$\{-3.5 \leq x \leq -2\} \cup \{2 \leq x \leq 3.5\}$	$\{-3.5 \leq y \leq -2\} \cup \{2 \leq y \leq 3.5\}$
37,79	$\{-2 \leq x \leq -0.5\} \cup \{0.5 \leq x \leq 2\}$	$\{-2 \leq y \leq -0.5\} \cup \{0.5 \leq y \leq 2\}$
113	$\{-0.5 \leq x \leq 0.5\}$	$\{-0.5 \leq y \leq 0.5\}$

The objective of the study on phase (φ) is to understand how the phase of the different frequency components of a signal is modified when the signal is filtered to some level with the Sign-Preserving filter. In this study, the amount of filtering is not defined by means of a W_{max} and n_{run} as it was done previously. The filter in this case is run until outputs \mathbf{b} and \mathbf{c}

of the same magnitude (defined as the Euclidean norm, see Eq. (18)) are produced. The reason for this (owing to the study of E_s) and details of how it is achieved will be explained further on.

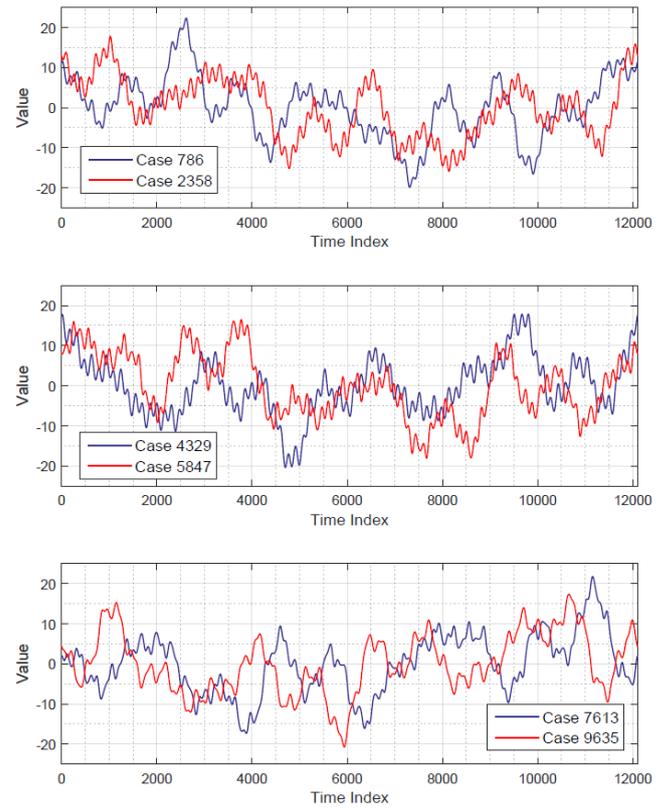


Figure 11. Examples of some randomly generated noisy-waves used for the study of φ and E_s

The results of the phase study are shown in Figures 12 and 13. For each of the 10,000 signals filtered, the difference in phase between the relevant harmonics (1, 4, 9, 11, etc) contained in the output signals and their analogous harmonic in the original signal \mathbf{a} is presented. Figure 12 shows the difference in phase ($\Delta\varphi$) or phase-shift of the frequency components of \mathbf{b} with respect to \mathbf{a} while Figure 13 presents the $\Delta\varphi$ between \mathbf{c} and \mathbf{a} . It should be noted that whilst the signals \mathbf{a} are composed exclusively of the harmonics listed in Table 2, the components \mathbf{b} and \mathbf{c} produced from the signal splitting have a richer frequency spectrum (i.e. have other harmonics). A direct comparison can only be carried out with those harmonics which have a counterpart in the original signal \mathbf{a} .

It can be seen that the Sign-Preserving filter does a fairly good job at conserving the phase of the frequency components. In the case of the low-frequency signal \mathbf{b} , shown in Figure 12, the $\Delta\varphi$ of all the slow harmonics (≤ 11), which are the most important in terms of magnitude, is no greater than 20° in either direction. In the 37th and 79th harmonics, differences of up to 180° are observed; however the majority of the results are concentrated in the $\pm 60^\circ$ band. The 113th harmonic is the only one whose $\Delta\varphi$ oscillates in the whole $\pm 180^\circ$ range; however, given that its magnitude in the \mathbf{b} signal is so small, this is not particularly relevant.

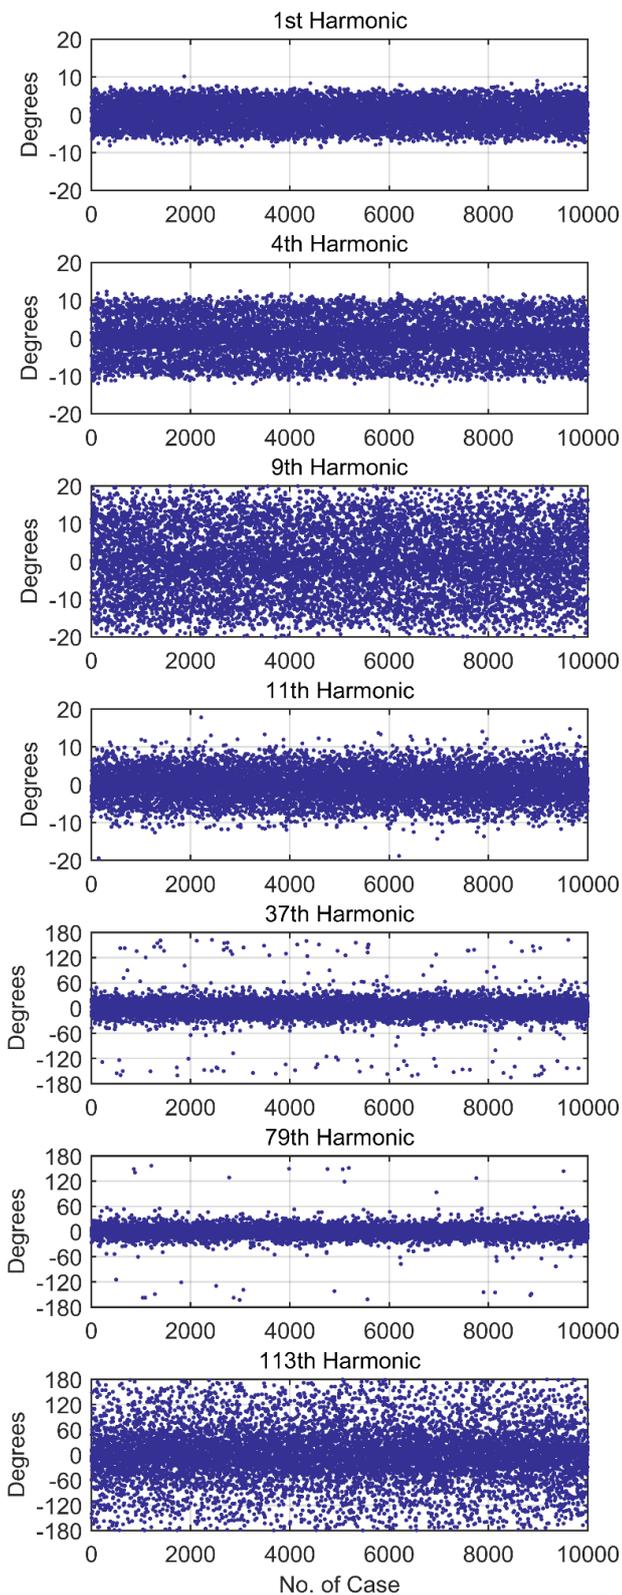


Figure 12. Phase difference between harmonics in **b** and their counterparts in **a** for the 10,000 random signals filtered

In the case of **c**—shown in Figure 13—the slow harmonics (≤ 11) maintain a $\Delta\varphi$ with respect to **a** well within a $\pm 20^\circ$ band, being the 9th harmonic the one that exhibits the most erratic behaviour. The 37th harmonic presents a difference of $\pm 5^\circ$ while the 79th harmonic is almost unaltered, having a shift within $\pm 1^\circ$. The results for the 113th harmonic are concentrated in the $\pm 10^\circ$ band, however a few outlying points can be distinguished.

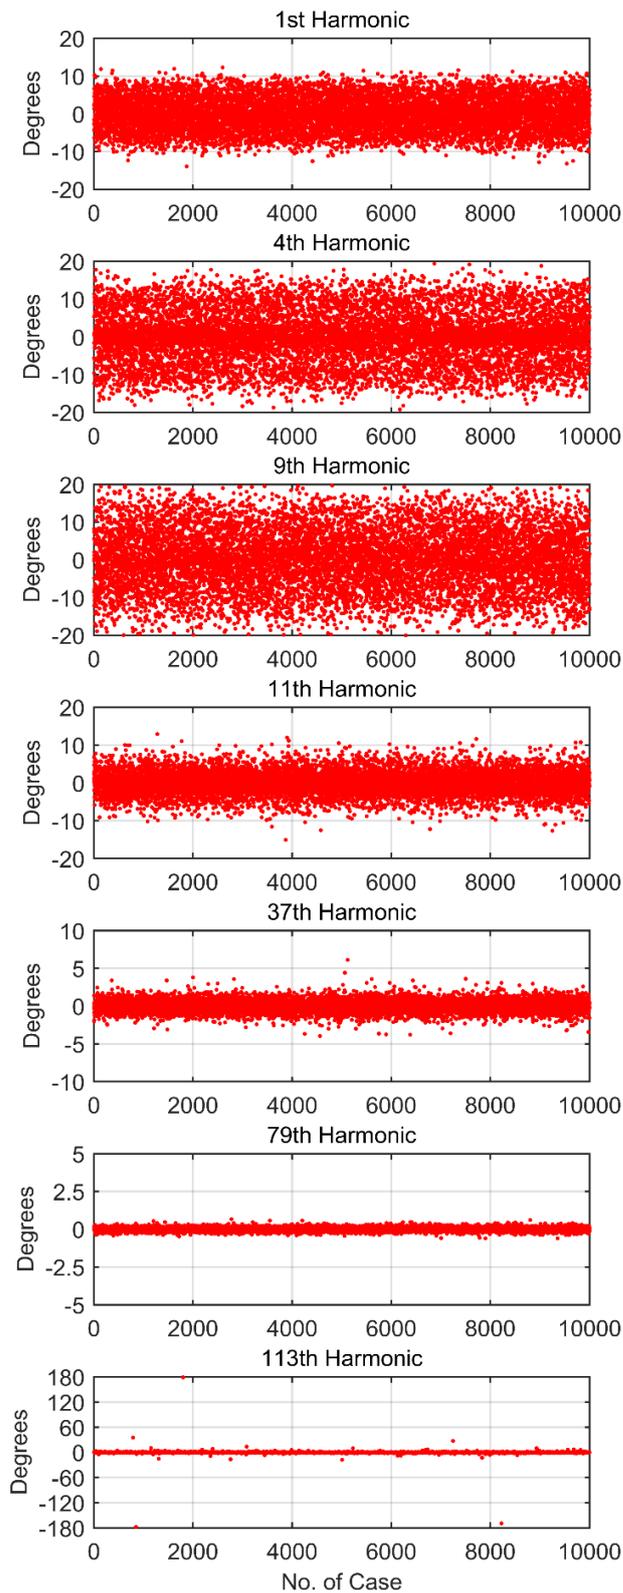


Figure 13. Phase difference between harmonics in **c** and their counterparts in **a** for the 10,000 random signals filtered

For the study on signal energy (E_s) a comparison between the Sign-Preserving filter and a more conventional Butterworth filter was carried out. The comparison between the two filters is not as straightforward as it seems because the Sign-Preserving filter, unlike the Butterworth, does not have a cut-off frequency as a control parameter; therefore generating comparable waves with the two methods is a challenge.

A solution to the above is to use the magnitude of the output signals as the stopping criterion. For each of the 10,000 randomly generated **a** waves, both filters will be used to

produce \mathbf{b} and \mathbf{c} components of equal magnitudes among themselves, thereby allowing a fair comparison.

For the Butterworth filter a loop was programmed so that the cut-off frequency could be varied until the value that produced a split of equal magnitudes was found. A low-pass filter is used to generate the \mathbf{b} signal while \mathbf{c} is simply assumed to be the remaining content (i.e. $\mathbf{c} = \mathbf{a} - \mathbf{b}$). In the case of the Sign-Preserving filter both control parameters are discrete and w_{max} is restricted to odd numbers; therefore obtaining output signals of equal magnitude by simply assigning a value to these two variables is virtually impossible. A loop that applies the filter repeatedly and allows a fine adjustment was constructed. In this loop the filter carries out as many runs as possible with a large W_{max} (to do the bulk of the work quickly), then it performs more runs with an increasingly smaller value of W_{max} (for a fine-tuning) until outputs of equal magnitudes are produced.

The signal energy (E_s) of a discrete signal \mathbf{f} is calculated as the sum of the squares of all the data points contained in the vector [14], as shown by Eq. (19). For each of the 10,000 cases explored, the signal energies of the original signal \mathbf{a} and of the output signals \mathbf{b} and \mathbf{c} produced by both filtering methods were calculated. A ratio of signal energies (Z) defined by Eq. (20) was also calculated.

$$E_s(\mathbf{f}) = \sum_{t=1}^L \mathbf{f}(t)^2 = (\text{norm}(\mathbf{f}))^2 \quad (19)$$

$$Z = \frac{E_s(\mathbf{b}) + E_s(\mathbf{c})}{E_s(\mathbf{a})} \quad (20)$$

Figure 14 shows a comparison of the ratios of signal energies (Z) obtained with the Butterworth filter and the Sign-Preserving filter. It can be seen that the Sign-Preserving filter consistently yields a lower Z than the Butterworth low/high pass filter. On average the ratio produced by the Sign-Preserving filter is 40% lower than the ratio obtained from the Butterworth filter.

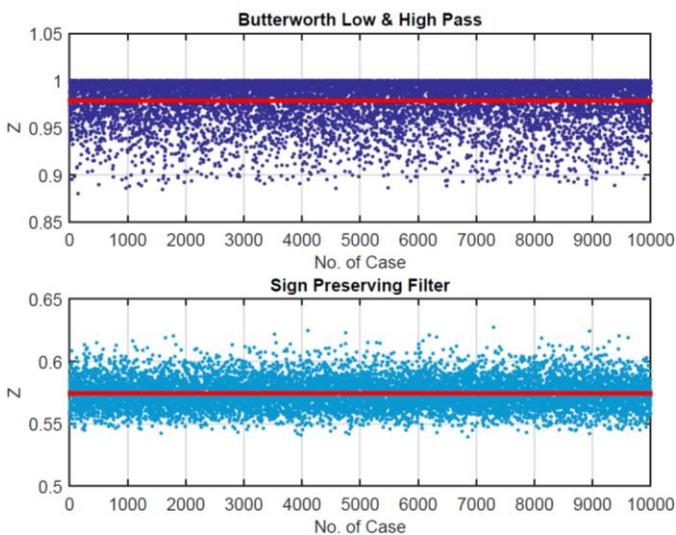


Figure 14. Comparison between the ratio of signal energies (Z) obtained through a Butterworth filter and the Sign-Preserving filter

Although the term signal energy is a measure of the size of the signal rather than an actual measure of energy, the two

concepts are very closely related. Considering an audio system for example, the E_s of a voltage signal represents the energy that would be dissipated by a 1Ω speaker [15]. The results presented in Figure 14 are of great importance because they indicate that in some applications, the Sign-Preserving filter potentially allows reproducing the effect of \mathbf{a} (since $\mathbf{b} + \mathbf{c} = \mathbf{a}$) but consumes less energy in doing so, since $E_s(\mathbf{b}) + E_s(\mathbf{c}) < E_s(\mathbf{a})$.

5 Examples of applications

In this section of the paper two examples of possible applications for the Sign-Preserving filter are presented. The filter is suitable for optimization problems which, similar to the examples presented, have the following conditions:

- 1) No counter-flow. The two output signals should never have different signs from each other.
- 2) Neither of the output signals should ever have a greater amplitude than the input signal.
- 3) The frequencies of the two output signals must be sufficiently different from each other but there is not a sharp cut-off frequency. In other words, one signal should be made mainly of low-frequency content while the other should mostly comprise high-frequency content.
- 4) An improvement in performance, cost or any other metric of interest can be attained by splitting the input signal into two (or more) signals and customizing the design of the physical components of the system that will handle (or will be driven by) them.

5.1 Optimization of an electric vehicle's battery pack

The battery pack of an electric vehicle (EV) can be cost-optimized by means of the Sign-Preserving filter. The power profile seen at the wheels of the vehicle (which is ultimately the load of the battery) can be split into two components, a predominantly low-frequency profile and a predominantly high frequency profile. Typically, the low-frequency component will have a large energy content and a reduced peak power, whereas the high-frequency component will have a small energy content but a comparatively large peak power.

Instead of having a single battery that handles the load in its entirety, the EV will have 2 batteries, one for each of the two frequency components (low and high) of the load. The two batteries will be based on different compositions or technologies. The optimization presupposes the possibility of manufacturing batteries of different specifications whose cost per unit energy and per unit power differ considerably from each other.

The “*High Capacity (HC)*” battery, used for the low-frequency part of the load, will have a much lower cost per unit of energy storage capacity ($\$/\text{kWh}$) while it will have a high cost per unit power ($\$/\text{kW}$). Conversely, the “*High Power (HP)*” battery, used for the high-frequency part of the load, will have a much lower cost per unit power but will have an increased cost per unit energy capacity. Therefore, the low-frequency battery will supply the bulk of the energy while the high-frequency battery will mainly function as a peak-shaver.

The High-Capacity battery of the proposed 2-battery hybrid system could be based for example, on a chemistry that has a

high specific energy such as Lithium Nickel Cobalt Aluminium-Oxide (NCA), Lithium Nickel Manganese Cobalt-Oxide (NMC) or even Sodium Nickel Chloride (Na/NiCl₂). The HP battery, on the other hand, could be based on a chemistry with a high specific power such as Lithium Iron-Phosphate (LFP) [16-18].

If a Fourier decomposition was used to split the signal, the problem of energy “counter-flow” would be encountered. In the context of an EV application having counter-flow implies one of two things: 1) one battery is supplying power to the wheels while at the same time the other battery is being recharged by the regenerative brakes or 2) the output of one battery is being used to charge the other. Needless to say, neither of those scenarios make sense, both batteries should be either charging or discharging at the same time. The case where one battery is idle while the other one is active is also possible because there is no work in opposite directions. Consequently, the Sign-Preserving filter is required.

Cardenas and Garvey [19] carried out, using the Sign Preserving filter, a cost-optimization of the battery pack of a Nissan Leaf. In the study, the load that the vehicle’s powertrain experiences when the car is subjected to a standard driving cycle [20] is used as the reference signal (a). The signal used for the study is created through a model that translates the speed profile of the driving cycle into an electric power profile [21]. Figure 15 shows the electric power profile of the Nissan Leaf when driving under the “EPA-LA92” cycle. During the positive periods the battery is supplying power to the electric motor whilst during the negative periods the battery is being recharged by the regenerative brakes.

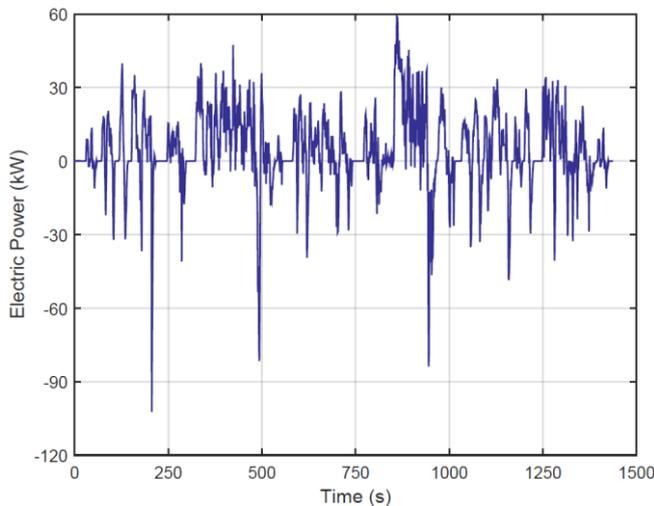


Figure 15. Electric power profile used for optimization

To carry out the optimization of the battery-pack a 2D combinatorial space was created with the two input parameters to the filter: W_{max} and n_{run} . The range for the variables was defined rather arbitrarily in the study; however the results show that the optimum is contained within the search space defined. Every combination of parameters ($W_{max} + n_{run}$) is a different filtering operation that produces two load profiles, one for each battery.

The original battery of the Nissan Leaf has a capacity of 24 kWh [22], which in the hybrid battery system is provided by the sum of the capacities of the two batteries: *High-Capacity* and *High-Power*. The fraction that either one of the batteries contribute to the total capacity of 24kWh of the vehicle is

called the “Capacity Ratio”. The capacity ratios of the High-Capacity (χ_{HC}) and High-Power (χ_{HP}) batteries can be calculated through Eqs. (21) and (22), respectively.

$$\chi_{HC} = \int_{t=0}^{t=T} b(t) dt \Big/ \int_{t=0}^{t=T} a(t) dt \quad (21)$$

$$\chi_{HP} = 1 - \chi_{HC} \quad (22)$$

The total system cost for each of the n combinations of *HC+HP* batteries contained in the two-dimensional search space (W_{max} and n_{run}) is calculated by means of Eqs.(23)-(25). The inputs for the equations are the capacity ratios and peak powers seen by the batteries in every different configuration.

$$cost_{HC} = 20 \cdot (24 \cdot E_{HC}) + 60 \cdot Peak_{HC} \quad (23)$$

$$cost_{HP} = 140 \cdot (24 \cdot E_{HP}) + 30 \cdot Peak_{HP} \quad (24)$$

$$cost_{total} = cost_{HC} + cost_{HP} \quad (25)$$

Figure 16 shows how the total system cost varies with respect to the two parameters of the filter. The highest overall costs observed are in the range between \$6700 and \$6800. The original battery pack of the Nissan Leaf has an estimated price of \$6805 [23].

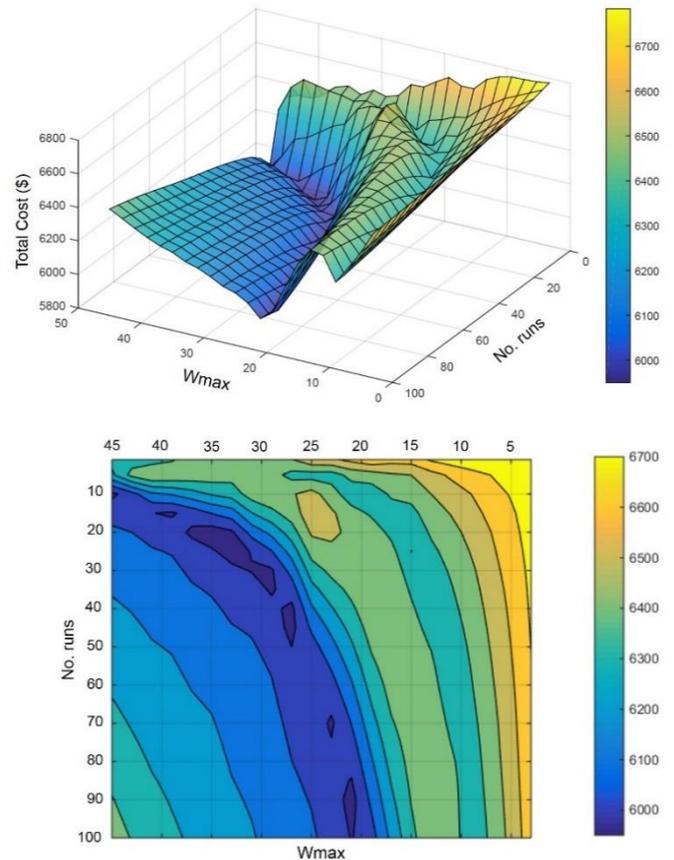


Figure 16. Variation of the total cost of the hybrid battery (HC + HP) with respect to the filter’s control parameters

The absolute minimum total cost of \$5939.3 is obtained with a combination of $W_{max} = 33$ and a $n_{run} = 22$. In the optimum configuration found (minimum cost) the High-Capacity battery has a capacity ratio (χ_{HC}) of 0.80, which is

equivalent to 19.23 kWh and it sees a maximum peak power of 37.18 kW. On the other hand, the High-Power battery has a capacity of 4.77 kWh and sees a much larger peak power of 88.55 kW. The cost of the optimum configuration found represents a 12.7% reduction in comparison to the cost of the original battery pack of the vehicle.

Figure 17 shows how the overall system cost varies with respect to the capacity ratio of the High-Capacity battery (χ_{HC}). The fact that there are many different combinations of filter parameters that yield the same capacity ratio becomes evident in the figure. It should be pointed out however, that even though the capacity ratio is the same, the shapes of the profiles generated are different. It is also clear from Figure 17 that for the same capacity ratio some combinations are more expensive than others.

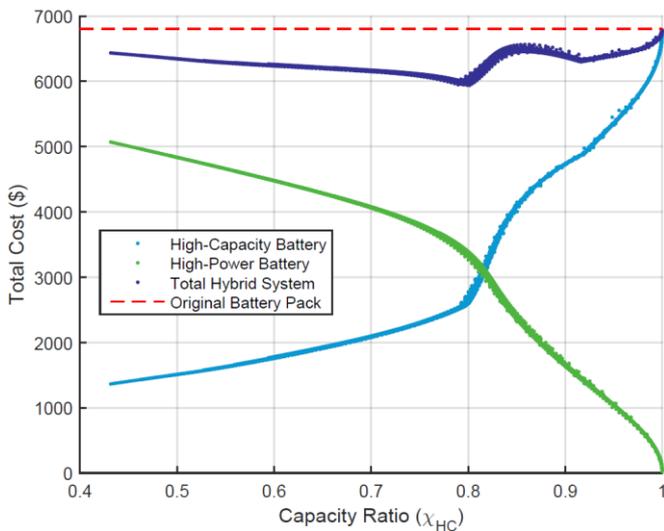


Figure 17. Relationship between the total cost of the hybrid battery and the capacity ratio of the High-Capacity battery (χ_{HC})

This case study demonstrates the usefulness of the “Sign-Preserving filter” in the context of electric vehicles. By replacing the battery pack of the vehicle with a hybrid system consisting of two batteries (one for bulk energy storage and one for peak-power) cost reductions of 12.7% (or more, depending on the load used as a reference) can be achieved.

In the study, the authors point out that the optimization approach used (based on the Sign-Preserving filter) is not limited to cars, but could be equally applied to other types of EVs. A hybridized battery would be particularly beneficial to vehicles that are likely to experience drastic or very frequent acceleration and braking, such as electric bikes and city buses.

5.2 Optimization of a packed bed thermal energy store

Another example of the use of the Sign-Preserving filter is the optimization of the thermal stores inside a compressed air energy storage system (CAES).

CAES is a large scale energy storage technology typically used for grid applications. During periods of excess electricity production (or low demand) a compressor draws power from the grid to compress air, which reaches high pressures (~ 70 Bar) and temperatures (~550 °C). The stream of high pressure air is circulated through a heat exchanger where it surrenders the heat of compression to a secondary stream of non-

pressurized air. The compressed air is then stored at high pressure but near ambient temperature in a dedicated pressure store. The secondary stream of air flows through a packed bed, in which the heat of compression is stored for a subsequent use.

During periods of high demand (or low electricity production) the CAES system operates in reverse to put energy back into the grid. The compressed air is withdrawn from the pressure store, heated up with the energy stored in the packed bed and expanded in a turbine. The mechanical work produced by the expansion of the air is used to drive an electrical generator. A more detailed explanation of the operation of CAES systems can be found in [24, 25]. Figure 18 shows a schematic of a packed bed working within a CAES system.

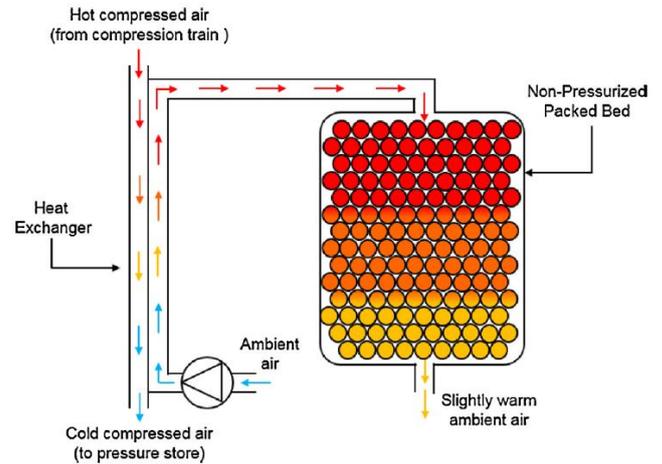


Figure 18. Schematic of a non-pressurized packed bed operating in charging mode within a CAES system.

There are different trade-offs involved in the design of a packed bed. For example, if small rocks are used the packed bed will have very good heat transfer characteristics due to a large surface area; however this entails increased pressure drops. Conversely if big rocks are used, pressure losses are minimized but the packed bed will suffer from poorer heat transfer capabilities due to the reduced surface area.

Another design trade-off is found in the choice of aspect ratio (α) for the container. The aspect ratio is defined as the ratio of the packed bed’s height with respect to its diameter. A short and wide design ($\alpha < 1$) has the advantage of a lower pumping load due to the reduced height and larger cross-sectional area of the container; however due to the reduced separation between the hot and cold ends of the store, the self-discharge exergy losses (heat flowing down the thermal front) are increased. A thin and tall packed bed ($\alpha > 1$) is the inverse case, it has much higher pressure drops due to an increased height and a reduced cross-section but has smaller self-discharge losses.

Cardenas et al. [26] carried out a techno-economic optimization of a utility-scale packed bed. In the study, the authors analyse comprehensively the aforementioned trade-offs. The work cycle (ω_1) used in the study is defined by Eq. (26); it is a power profile with a peak power of 10 MW and a period of 1 day (86400 sec.)

$$\omega_1 = 10 \cdot \sin\left(\frac{2 \cdot t \cdot \pi}{86400}\right) \quad (26)$$

The work cycle ω_1 has 1 charging period (between the 0th and 12th hour) and 1 discharging period (between the 12th and 24th hour). The work cycle has an energy content (Q) of 76.4 MWh_{th}, calculated via Eq.(27), which are supplied to the packed bed by the secondary stream of non-pressurized air at 550°C.

$$Q = \int_{t=0}^{t=43200} \omega_1 \cdot dt \quad (27)$$

The mass of rock required to store the energy content of the work cycle (Q) is determined based the specific heat capacity of the rocks and the temperature delta that the packed bed will experience, which in a CAES system goes from ambient to 550 °C. The packed bed of the study contains 850x10³ kg of rock.

The parametric study carried out explored several different designs of packed beds (combination of aspect ratio and size of rocks) to find the optimum configuration from a techno-economic standpoint, i.e. a design that achieves the highest efficiency with the minimum cost.

Expanding on the above, Cardenas et al. [27] presented a brief analysis of the effect that the frequency of the work cycle has on the performance of a packed bed. For the study, different work cycles were created by reducing the period of ω_1 to a half (ω_2), a fourth (ω_4), an eight (ω_8), etc. as Eqs. (28) and (29) show. In order for the packed beds of the different work cycles to be comparable to one another, their size (based on the energy content of the cycle) was kept constant. To maintain the 76.4 MWh_{th} of energy storage capacity, the amplitude (peak power) of the profiles is modified proportionally to the frequency; in other words, if the frequency is doubled the peak power of the cycle is doubled too.

$$\omega_2 = 20 \cdot \sin\left(\frac{2 \cdot t \cdot \pi}{43200}\right) \quad (28)$$

$$\omega_4 = 40 \cdot \sin\left(\frac{2 \cdot t \cdot \pi}{21600}\right) \quad (29)$$

For clarity, ω_2 refers to a 24 hour work-cycle with 2 charging periods and 2 discharging periods (of equal duration) while ω_4 refers to a 24 hour work-cycle with 4 charging periods and 4 discharging periods. Figure 19 shows the work cycles used for the frequency study.

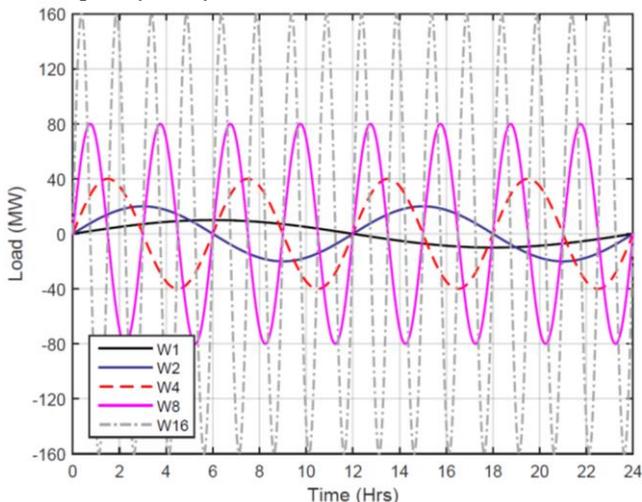


Figure 19. Different 76.4MWh work-cycles used for frequency study

The parametric analysis carried out with ω_1 in ref. [26] was repeated with the different work cycles ω_2 , ω_4 , etc. For every one of the load profiles several different values of α and sizes of rocks were evaluated to find the optimum design in terms of performance (lowest exergy losses).

Figure 20 shows the behaviour of the exergy losses per cycle (normalized with respect to the exergy input) for a packed bed working under ω_2 (top) and ω_8 (bottom) as the aspect ratio and rock size are varied. It can be seen that regardless of the design parameters, the exergy losses per cycle increase as the frequency of the work-cycle increases. The optimum design found for the ω_1 work cycle had losses of 1.75% whereas the optimum designs for work-cycles ω_2 and ω_8 have exergy losses of 2.31% and 3.46%, respectively

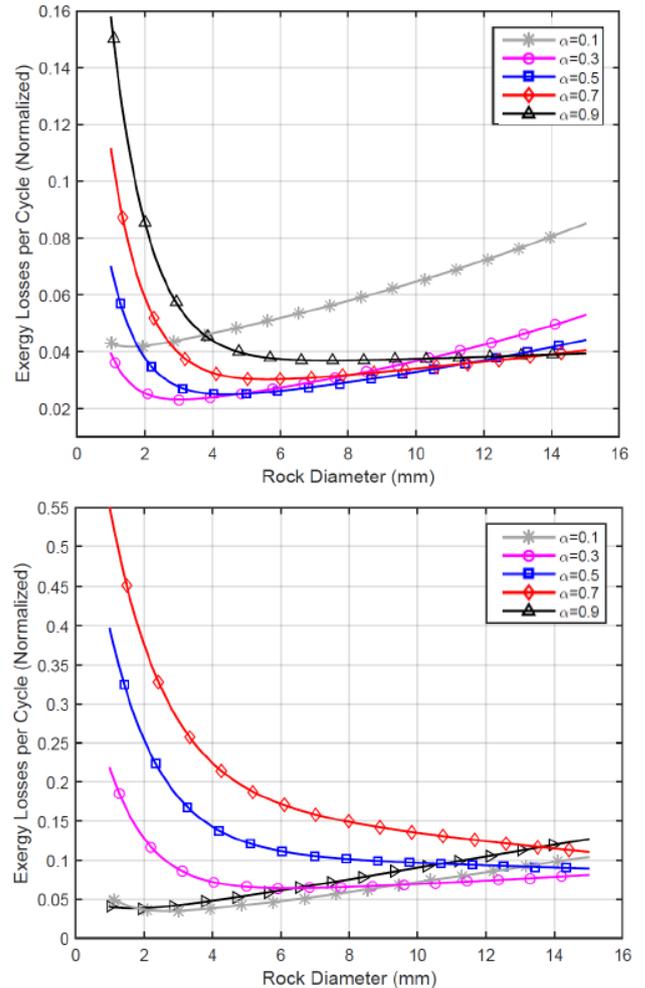


Figure 20. Exergy losses per cycle for different designs of packed beds working under a ω_2 (top) and a ω_8 (bottom)

Figure 21 shows a more generalized view of the behaviour of the exergy losses for different working cycles. All the designs shown in the figure consider the optimum size of rocks for the specific value of α . It can be seen that the exergy losses per cycle increase as the frequency of the work-cycle increases. This is largely attributable to the fact that higher frequency work-cycles have a higher “peak-power to energy” ratio, which translates directly into much higher pressure losses. The optimization algorithm will try to compensate this by shifting the design of the packed bed towards a smaller α ; however the extent of the adjustment is limited by the other mechanisms of exergy loss (such as heat transfer and self-discharge losses) which tend to increase as α reduces.

On the other hand, for the case of the lower frequency work-cycles, the design of the packed bed will tend to larger values of α to try to reduce heat transfer and self-discharge losses since the lower “peak-power to energy” ratios of these cycles do not entail particularly high pressure drops.

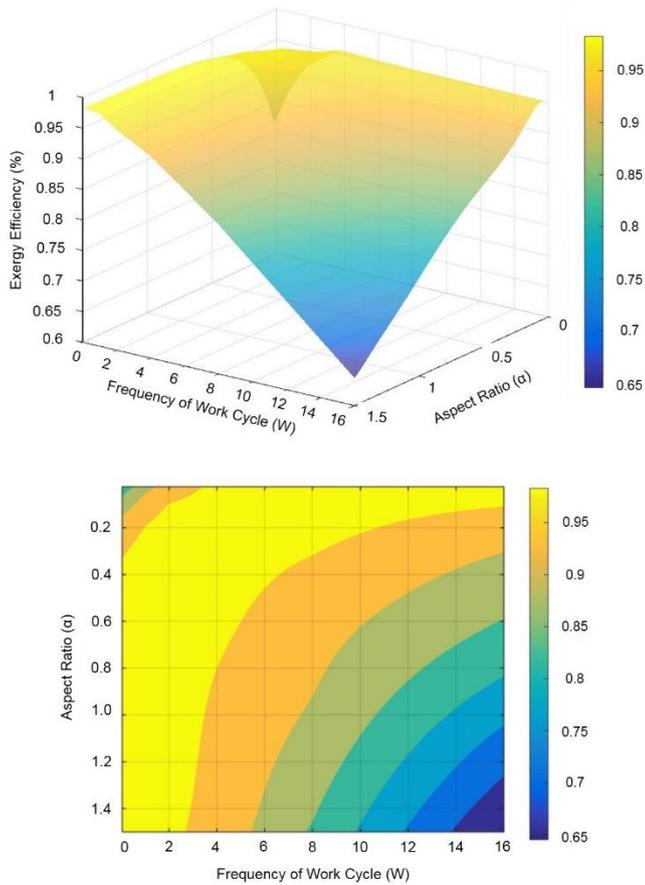


Figure 21. Effect of the frequency of the work-cycle on the performance of the packed bed.

Another interesting detail to note from Figure 20 is that the rock size for the optimum design becomes increasingly smaller as the frequency of the work cycle increases. The increased air flow rates (due to higher powers) seen in the faster work cycles demand a better heat transfer capability from the packed bed. The specific surface area of the packed bed can be increased if smaller rocks are used, which allows a more effective heat exchange between the storage material and the stream of air; however smaller rocks will result in increased pressure drops.

Table 3. Optimum packed bed designs obtained for the different work-cycles studied.

Work Cycle	Aspect Ratio (α)	Rock Diameter (mm)	Normalized Exergy Losses	Exergy Efficiency (%)
$\omega_{0.5}$	1.2	4.36	0.0154	98.46
ω_1	0.6	3.72	0.0174	98.26
ω_2	0.3	3.06	0.0231	97.69
ω_4	0.2	3.03	0.0289	97.11
ω_8	0.1	2.67	0.0346	96.54
ω_{16}	0.05	2.26	0.0423	95.77

Table 3 summarizes the design and performance parameters of the optimum designs found for each of the 6 work cycles studied. As it can be drawn from the results presented, there is a design trade-off involving the frequency of the work-cycle. If a packed bed is intended for operating under a high-

frequency work cycle, small rocks will have to be used in order to sustain the rapid heat transfer required; although doing so will result in a higher pumping load. Conversely, if a packed bed will operate under a low-frequency work cycle, a fast-response capability is not necessary and larger rocks can be used to minimize pressure losses.

The work cycle that a CAES system is likely to experience in real life applications is not a simple sinusoidal wave. It is a profile that is determined by electricity production and demand and has considerable oscillations (noise) in it. The trade-off discussed above leads to the idea that it is possible to achieve a reduction in the total exergy losses of a packed bed (either working in a CAES system or other type of application) if instead of servicing a load with a single packed bed, the load is split based on frequency ranges into two (or more) smaller profiles and an independent packed bed is assigned to each one of them.

If a Fourier decomposition was used to split the signal, the problem of energy “counter-flow” would be encountered. In the context of a CAES system the counter-flow implies that a single reversible compressor/expander machine cannot be used because at some points one packed bed would be in charge mode (i.e. compressor working) while at the same time the other would be discharging (i.e. expander working) therefore two separate power-conversion machines would be required.

Furthermore, when counter-flow occurs, the output of one packed bed is used to counteract the input of the other one to create the desired overall effect which is to replicate the original profile. Because of this, one of the two packed beds will see a greater power than what the single packed bed for the original profile would see at the same time (Fig. 1a illustrates this), which calls for a power-conversion machine that is rated for a higher power and is thus more expensive. For the reasons above, the Sign-Preserving filter is required to carry out the load splitting.

The expected reduction in exergy losses can be achieved because by splitting the load into multiple components, the design of each one of the individual packed beds can be optimized better for the narrower range of frequencies to which they will be subjected. The low-frequency component of the system’s load profile, which will typically have the largest energy content, will be handled by a packed bed with large rocks which will keep pressure losses to a minimum. On the other hand, the higher frequency component(s) of the load, which hold a much smaller energy content, will be handled by a packed bed with rather small rocks. A penalty in pressure drops will be paid in favour of improved heat transfer characteristics; however the total amount of exergy that will pass through the high-frequency packed bed will be relatively small.

Different combinations of the filter’s control variables W_{max} and n_{run} need to be explored to find the optimum split of the power profile. Each combination will produce two load profiles (**b** and **c**) which will be handled by two different packed beds. The design of each packed bed will be optimized for its working load, following the procedure described in [26]. The optimum signal split is the one that produces the lowest total exergy losses.

The work on the optimization of a packed bed based on “load splitting” is currently being carried out by the authors of this paper. The concept and rationale behind this application have been introduced in [27] and [28]. The results of the optimization of a packed bed by means of the Sign-Preserving filter will be made available in a subsequent publication.

6 Concluding remarks

There are optimization problems where a reduction in cost or an improvement in performance can be attained if the driving signal of the system is separated into two or more components, because this allows customizing the design of the physical equipment for a narrower range of frequencies and thus making a better use of its properties.

Some of those problems, such as the hybridization of the battery pack of an electric vehicle or the parallelization of a packed bed thermal store, have very specific signal-splitting requirements that conventional signal processing tools cannot meet. These include:

- The two output signals should never have different signs from each other.
- Neither of the output signals should ever have a greater amplitude than the input signal.
- One signal should comprise mainly low-frequency content while the other should comprise mostly high-frequency content.

A novel Sign-Preserving filter for discrete signal decomposition has been presented in this paper and the mechanics of its operation have been thoroughly discussed. The filter developed splits an original signal **a** into two components: a mostly low-frequency signal **b** and a predominantly high-frequency signal **c**. The sum of **b** + **c** replicates **a** perfectly. The key feature of the filter is that the signs of **b** and **c** are equal to the sign of the original signal **a** at every time *t*.

The algorithm followed by the Sign-Preserving filter has been explained in depth. The filter produces 2 outputs, although if a 3-way split is required the filter can be reapplied to the low-frequency output **b**. It should be noted that the filter does not operate in real-time, however a quasi-real-time operation could be achieved if forecast data is available for the particular application. The operation of the filter is rather simple, it is based on passing wavelets (similar to a Hann window) of different widths through the signal length subtracting some amount of “non-smoothness” from it with every pass which is then stored in the high-frequency output signal (**c**). A set of checks are in place to ensure that no sign-change occurs in neither of the signals (**b** and **c**). The filter will perform passes along the signal with all the wavelets contained between 3 and the maximum wavelet width defined by the user. A run of the filter consists of two passes with every wavelet contained in the aforementioned range. Several runs can be carried out, which will produce a smoother low-frequency signal **b**.

The operation of the Sign Preserving filter has been demonstrated with a manufactured noisy-wave (whose equation of provided for reproducibility) and with square waves of different lengths. Not a single case was found where

the sign-preservation condition is not met. The paper also presents a study on phase, using 10,000 randomly generated noisy-waves. The study shows that, as a consequence of its sign-preserving characteristic, the filter also preserves the phase of the different frequency components. All the slow harmonics in the low-frequency signal **b** produced (in each of the 10,000 cases analysed) have a difference of $\pm 20^\circ$ with respect to their counterparts in the original signal, while the fast harmonics contained in the high-frequency signal **c** have a difference of only $\pm 10^\circ$.

With the same 10,000 random waves a study on signal energy was carried out. A parameter *Z* was defined as the ratio of the sum of the signal energies of **b** and **c** with respect to the signal energy of the original signal **a**. The *Z*-values produced by a low + high pass Butterworth filter range between 0.88 and 1 whilst the values for *Z* generated by the Sign-Preserving filter range between 0.53 and 0.63. These results indicate that the Sign-Preserving filter is capable of reproducing the effect of a signal but consumes less energy in doing so.

Furthermore, the paper presented two case studies of the use of the Sign Preserving filter for real life optimization problems where “counter-flow” needs to be avoided. One is the cost reduction of an electric vehicle’s battery pack by replacing it with two batteries: A “High-Capacity” one with a low cost per unit of energy storage capacity (\$/kWh) and a “High-Power” one with a low cost per unit power (\$/kW). The power profile of the vehicle is split into two signals by means of the Sign Preserving filter, the low-frequency part of the load is handled by the High-Capacity battery while the High-Power battery takes care of the high-frequency part of the duty. The case study showed that a cost reduction of 12.7% or more (depending on the vehicle, driving profile and batteries used) can be achieved through this optimization approach based on load-splitting.

The second case study is concerned with the packed beds of rock (thermal stores) inside a compressed air energy storage system. Although conclusive results have not been achieved yet, a strong case for the use of the Sign Preserving filter in this application is made. If a packed bed is intended for operating under a high-frequency work cycle, a small rocks are needed to sustain the rapid heat transfer required; however they entail higher pressure drops. Conversely, if a packed bed will operate under a low-frequency work cycle, larger rocks can be used because a fast-response capability is not necessary and pressure losses are minimized.

The work cycle that a packed bed operating inside a CAES system experiences is a profile with considerable oscillations in it, i.e. mix of a wide range of frequencies. Based on the results presented, it seems very possible that an improvement in the efficiency of a thermal energy store can be achieved if the load is split into two (or more) smaller profiles based on frequency ranges and each profile is handled by an independent packed bed whose design has been customized for it.

Acknowledgments

This research work has been carried out within the framework of the following two projects funded by the UK’s Engineering and Physical Sciences Research Council (EPSRC): 1) “Next

Generation Grid Scale Thermal Energy Storage Technologies” (EP/L014211/1) and 2) “Multi-scale Analysis for Facilities for Energy Storage” (EP/N032888/1).

References

[1] I.F. Apolinario, P.S.R. Diniz. Chapter 1: Introduction to Signal Processing Theory in S. Theodoridis (Ed.), Academic Press Library in Signal Processing, 2014, pp.3-28. Cambridge, MA. Academic Press.

[2] M. Ferrer, M.d.Diego, G.Piñero, A.Gonzalez. Active noise control over adaptive distributed networks. *Signal Processing* 2015;107: 82–95

[3] O.J. Kwon, S. Choi, D.Jang. H.S.Pang. All-in-focus imaging using average filter-based relative focus measure. *Digital Signal Processing* 2017; 60:200–210

[4] A. Chalmers, K. Debattista. HDR video past, present and future: A perspective. *Signal Processing: Image Communication* 2017; 54: 49–55

[5] S.Y. Tsui, C.S. Liu, C.W Lin. Modified maternal ECG cancellation for portable fetal heart rate monitor. *Biomedical Signal Processing and Control* 2017; 32:76–81

[6] P.K. Jain, A.K. Tiwari. An adaptive thresholding method for the wavelet based denoising of phonocardiogram signal. *Biomedical Signal Processing and Control* 2017; 38:388-399.

[7] A.C.To, J.R. Moore, S.D.Glaser. Wavelet denoising techniques with applications to experimental geophysical data. *Signal Processing* 2009; 89:144–160.

[8] M.T. Özgen, E.E. Kuruoglu, D. Herranz. Astrophysical image separation by blind time–frequency source separation method. *Digital Signal Processing* 2009; 19:360-369.

[9] C.C. Tseng, S.L. Lee. Closed-form designs of digital fractional order Butterworth filters using discrete transforms. *Signal Processing* 2017; 137:80–9

[10] J. Piskorowski. A new concept of phase-compensated continuous-time Chebyshev filters. *Signal Processing* 2008; 88: 437–447.

[11] D. Schlichthärle. Digital filters: Basics and design. Chapter 2: Analog filters. 2011, pp 13-50. Berlin, DE. Springer-Verlag.

[12] D. David. Sound system engineering. Chapter 12: Loudspeakers and Loudspeaker Arrays. 2006, pp. 286-300. Burlington, MA, U.S. Focal Press.

[13] Ç. Candan, H. Inan. A unified framework for derivation and implementation of Savitzky–Golay filters. *Signal Processing* 2014; 104:203–211

[14] B.P. Lathi, R.A. Green. Essentials of Digital Signal Processing. 2014, pp. 21-23. NY, U.S. Cambridge University Press.

[15] G.E. Carlson. Signal and Linear System Analysis. 1993. Boston, MA, U.S. Houghton Mifflin Company

[16] G. Zubi, R. Dufo-Lopez, M. Carvalho, G. Pasaoglu. The lithium-ion battery: State of the art and future perspectives. *Renewable and Sustainable Energy Reviews*, 2018; 89:292-308.

[17] A.M. Andwari, A. Pesiridis, S. Rajoo, R. Martinez-Botas, V. Esfahanian. A review of battery electric vehicle technology and readiness levels. *Renewable and Sustainable Energy Reviews*, 2017; 78:414-430.

[18] A. Satki, J.J. Michalek, E.R.H. Fuchs, J.F. Whitacre. A techno-economic analysis and optimization of Li-ion batteries

[19] B. Cárdenas, S.D. Garvey. Load optimization for reducing the cost of an electric vehicle’s battery pack. *Journal of Energy Storage* (2018) In Press.

[20] US Environmental Protection Agency. Dynamometer Drive Schedules. Available online at:

<https://www.epa.gov/vehicle-and-fuel-emissionstesting/dynamometer-drive-schedules>

[21] C. Fiori, K. Ahn, H. Rakha. Power-based electric vehicle consumption model: Model development and validation. *Applied Energy*. 2016; 168:257-268.

[22] Automotive Energy Supply Corporation (2018). Cell, module and pack for EV applications. Available online at: http://www.eco-aesc-lb.com/en/product/liion_ev/

[23] J. Burn (8 December 2014) Nissan leaf battery replacement to cost £4920. *Auto Express UK*. Retrieved from: <http://www.autoexpress.co.uk/nissan/89694/nissan-leafbattery-replacement-to-cost-4920>

[24] E. Barbour, D. Mignard, Y.Ding, Y.Li. Adiabatic Compressed Air Energy Storage with packed bed thermal energy storage. *Applied Energy*. 2015; 155: 804-815

[25] B. Cárdenas, A.J. Pimm, B. Kantharaj, M. C. Simpson, J.A. Garvey, S.D. Garvey. Lowering the cost of large-scale energy storage: High temperature adiabatic compressed air energy storage. *Propulsion and Power Research*. 2017; 6(2): 126-133.

[26] B. Cárdenas, T.R. Davenne, J.Wang, Y.Ding, Y.Jin, H.Chen, Y.Wu, S.D. Garvey. Techno-economic optimization of a packed-bed for utility-scale energy storage. *Applied Thermal Engineering*. 2018. In Press

[27] B. Cárdenas, T.R. Davenne, J.P. Rouse, S.D. Garvey. Effect of design parameters on the exergy efficiency of a utility-scale packed bed. *Journal of Energy Storage* 2018; 18:267-284.

[28] B. Cárdenas (2016, December) Multiple Thermal Stores for Reducing Exergy Losses in a CAES System. Presented at the UK Energy Storage Conference 2016 held in Birmingham, UK.