

Advanced and versatile signal conditioning for GNSS receivers using the high-rate DFT-based data manipulator (HDDM)

Johannes Rossouw van der Merwe  | Fabio Garzia  | Alexander Rügamer  | Wolfgang Felber 

Satellite Based Positioning Systems
Department, Fraunhofer IIS, Germany

Correspondence

Johannes Rossouw van der Merwe, Satellite Based Positioning Systems Department, Fraunhofer IIS, Germany.
Email: johannes.roussouw.vandermerwe@iis.fraunhofer.de

Present address

Nordostpark 84, 90411 Nuremberg,
Germany

Abstract

Proper signal conditioning is crucial for reliable Global Navigation Satellite System (GNSS) navigation. Signal conditioning includes correcting receiver front-end distortions, shaping noise, removing interferences, and altering the received signal. The high-rate DFT-based data manipulator (HDDM) is a versatile signal processing architecture based on the Discrete Fourier Transform (DFT). In this article, the HDDM is theoretically modeled, analyzed, and evaluated. Some applications are presented, including interference mitigation, spectrum reconstruction, overlay signal design, altering signal modulations, and signal equalization. Additionally, the complexity of one concrete hardware implementation is investigated. The HDDM has shown excellent results with interference mitigation, and it can simultaneously achieve other signal conditioning tasks. Processing several tasks with the same nested architecture provides a significant processing benefit compared to discrete processing architectures. It emphasizes the benefits of a single architecture to simultaneously address several signal condition tasks, as opposed to separate structures requiring significantly more processing or poorer synergy.

KEYWORDS

Discrete Fourier Transform (DFT), equalization, Global Navigation Satellite System (GNSS), high-rate DFT-based data manipulator (HDDM), interference mitigation, signal conditioning, spectrum compression

1 | INTRODUCTION

The antenna, the Radio-Frequency Front-End (RFFE), and digital signal conditioning of a Global Navigation Satellite System (GNSS) receiver significantly impact its position, navigation, and timing (PNT) performance. Although these stages are not directly involved in GNSS processing, they determine which signals are received, the received

signal gain (Popugaev et al., 2007), the suppression of unwanted signals (O'Brien & Gupta, 2010), limit distortions (Phelts et al., 2009), and shape the noise (Bruton, 1997). Therefore, these initial stages in a receiver processing chain should be designed and developed carefully.

Low-performance components (usually associated with low-cost designs) degrade the navigation quality to be barely functional in the worst-case scenario. In other

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2021 The Authors. *NAVIGATION* published by Wiley Periodicals LLC on behalf of Institute of Navigation.

words: garbage in, garbage out (GIGO) (Smith, 1994). However, the problem is that high-quality antennas, RFFE, and digital platforms for signal conditioning are typically expensive to develop and manufacture. Therefore, there is an inherent trade-off between high performance and low size, weight, power, and cost (SWAP-C) for system design.

Signal conditioning refers to the receiver pre-processing tasks, where noise is suppressed, interference is removed, and signal conditioning aims to provide a receiver with the best possible signal quality before the primary GNSS processing starts (i.e., acquisition and tracking for GNSS receivers).

This article presents the high-rate DFT-based data manipulator (HDDM) as a flexible signal conditioning architecture. The focus is on wideband GNSS receivers with a bandwidth exceeding 40 MHz. The HDDM was initially proposed as an interference mitigation technique (van der Merwe et al., 2020a) but has since proven useful for other signal conditioning tasks (van der Merwe et al., 2020b).

The HDDM deploys maximally overlapping Discrete Fourier Transforms (DFTs). It isolates the signal spectrum into multiple sub-bands without down-sampling the signal. Therefore, it allows efficient time-frequency mitigation of interferences, restructuring the signal's spectrum, overlaying different frequency bands, and equalizing the received signal.

The contribution of this paper is two-fold: First, it provides a theoretical foundation of the method to address shortcomings in the literature—it includes mathematical modeling, comparison to similar techniques, and complexity analysis; second, it showcases multiple use cases for this technique. The HDDM is moderately processing-intensive when it only fulfills a single task. However, if it addresses several signal conditioning functions simultaneously, it significantly saves on processing resources.

Therefore, the authors suggest a multi-conditioning approach to take advantage of this architecture. Discrete signal conditioning blocks (e.g., separating interference mitigation and equalization) excel at single tasks, but often do not synergize well with each other or require duplicate processing that increases SWAP-C of the receiver. The core concept of the HDDM is to overcome this limitation by removing duplicate operations and streamlining the signal conditioning architecture.

This paper is structured as follows: Section 2 presents a theoretical model of the HDDM, analyzes the expected signal processing capabilities, and compares it to other methods; Section 3 discusses possible hardware implementations and practical considerations of the algorithm; Sections 4 to 6 present three applications of the HDDM for improved GNSS signal conditioning; interference monitoring is showcased in Section 4; Section 5 shows how

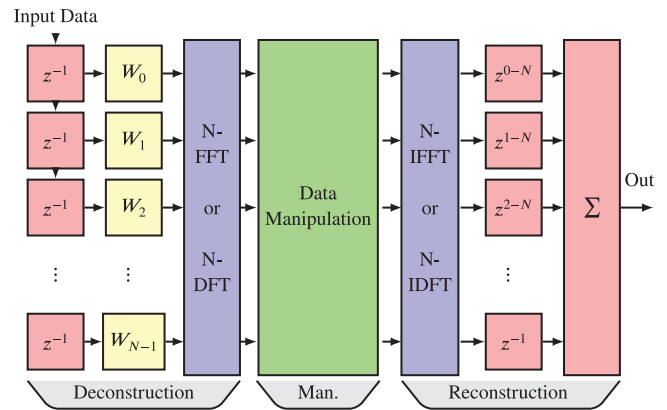


FIGURE 1 HDDM block diagram [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

the GNSS spectrum can be restructured; and equalization to correct RFFE distortions is demonstrated in Section 6; Section 7 provides a general discussion and outlook; and finally, Section 8 draws some conclusions.

2 | THE HDDM ALGORITHM

The HDDM algorithm was first introduced as a software implementation for interference mitigation (van der Merwe et al., 2020a) and later expanded to hardware-based interference mitigation (Garzia et al., 2020a, 2020b). Since then, other applications for the HDDM have also been investigated and demonstrated, such as spectrum compression and overlay-receiver designs (van der Merwe et al., 2020b). However, formal theoretical analysis is still missing. Figure 1 shows the generalized block diagram for the HDDM.

This section presents the theory behind the HDDM. It includes an impulse-response model for each channel and the computational complexity of the design. The HDDM is compared to a generalized channelizer architecture to highlight the complexity and performance differences and demonstrate its novelty.

2.1 | Signal conditioning architectures

2.1.1 | Single-rate channelizer architectures

This section derives the theoretical throughput for the HDDM. Figure 2 shows the block diagram of a channelizer receiver and the HDDM.

First, the input data stream is placed in a shift register to parallelize the data. A window function then temporally weighs the input samples. The windowed signal $x_w[n, m]$ for the n -th sample and the m -th channel can be written

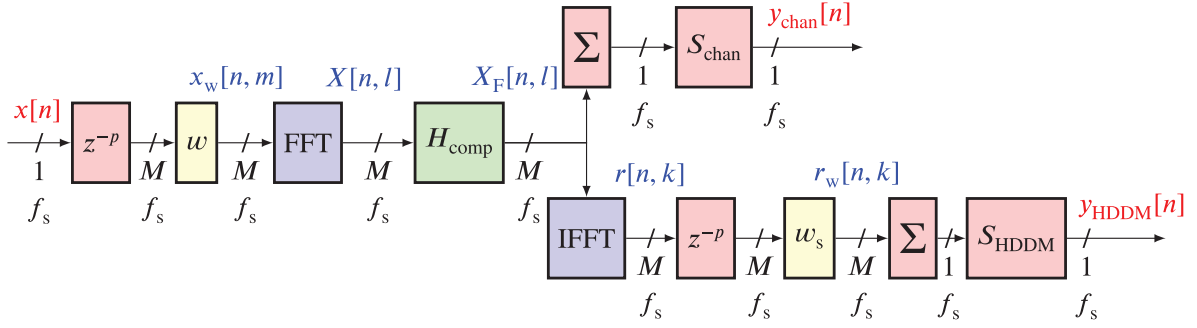


FIGURE 2 Block diagram for a channelizer receiver $y_{\text{chan}}[n]$ and the HDDDM $y_{\text{HDDM}}[n]$ [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

as:

$$x_w[n, m] = x[n - m] \cdot w[m]; 0 \leq m \leq M - 1 \quad (1)$$

where $x[n]$ is the digital input signal, and $w[m]$ is the window function. A DFT—or the more processing efficient Fast Fourier Transform (FFT)—transforms the windowed data to the frequency domain. The resultant frequency channelized data $X[n, l]$ is:

$$X[n, l] = \mathcal{F}\{x_w[n, m]\} \quad (2)$$

$$= \sum_{m=0}^{M-1} x_w[n, m] \cdot e^{\frac{-2\pi j}{M}lm} \quad (3)$$

where $\mathcal{F}\{\cdot\}$ denotes the FFT operation. The change of the channel index variable from m to l emphasizes the domain transformation.

Once the data is in the frequency domain, it can be easily manipulated, depending on the signal conditioning goal. A processing function $H_{\text{comp}}(\cdot)$ alters the channelized data to obtain the spectrum manipulated data $X_F[n, l]$:

$$X_F[n, l] = H_{\text{comp}}(X[n, l]) \quad (4)$$

The function $H_{\text{comp}}(\cdot)$ can be non-linear, time-variant, or interleave between the channels. However, in the simplest case, the function is assumed to be a linear and time-invariant scaling function:

$$X_F[n, l] = X[n, l] \cdot H[l] \quad (5)$$

$$= H[l] \sum_{m=0}^{M-1} x[n - m] \cdot w[m] \cdot e^{\frac{-2\pi j}{M}lm} \quad (6)$$

where the filtering function $H[l]$ consists of complex coefficients for spectrum shaping and filtering. This

assumption is retained for the remainder of this section. More complicated functions are considered in later sections.

The next step is to recombine the spectrum manipulated data $X_F[n, l]$ into a single processing stream. The recombination differentiates the HDDDM method from a channelizer architecture.

The FFT for every sample offset results in multiple, parallel, band-limited channels. Therefore, the most straightforward method is to combine these channels:

$$y_{\text{chan}}[n] = \frac{1}{S_{\text{chan}}} \sum_{l=0}^{M-1} X_F[n, l] \quad (7)$$

This output $y_{\text{chan}}[n]$ is a generic channelizer. A fixed gain S_{chan} compensates the window gain and ensures unitary throughput gain of the system:

$$S_{\text{chan}} = \sum_{m=0}^{M-1} w[m] \quad (8)$$

To obtain the transfer function for the system, let us assume:

$$H[l] = A \cdot \delta[l - l_{\text{ch}}] \quad (9)$$

where A is a constant gain, l_{ch} represents the channel selected, and $\delta(\cdot)$ is the Kronecker delta function defined as:

$$\delta[n] = \begin{cases} 1 & \text{if } n = 0 \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

Through substitution, it can be proven that:

$$y_{\text{chan}}[n] = \frac{A}{S_{\text{chan}}} \sum_{m=0}^{M-1} x[n - m] \cdot w[m] \cdot e^{\frac{-2\pi j}{M}l_{\text{ch}}m} \quad (11)$$

$$= x[n] * \left(\frac{A}{S_{\text{chan}}} \cdot w[n] \cdot e^{-\frac{2\pi j}{M} l_{\text{ch}} n} \right) \quad (12)$$

$$= x[n] * h_{\text{ch}}^{(\text{chan})}[n, l_{\text{ch}}] \quad (13)$$

where $*$ denotes the convolution operation, and $h_{\text{ch}}^{(\text{chan})}[n, l_{\text{ch}}]$ is the equivalent band-pass filter. It proves that each channel is a band-limited filtered channel and that no Inverse Fast Fourier Transform (IFFT) is required. However, the HDDM uses an IFFT to reconstruct the signal. The signal should be further analyzed to understand the necessity for the IFFT. The IFFT data is defined as:

$$r[n, k] = \mathcal{F}^{-1}\{X_{\text{F}}[n, l]\} \quad (14)$$

$$= \frac{1}{M} \sum_{l=0}^{M-1} X_{\text{F}}[n, l] \cdot e^{\frac{2\pi j}{M} lk} \quad (15)$$

$$= \frac{1}{M} \sum_{l=0}^{M-1} H[l] \sum_{m=0}^{M-1} x[n-m] \cdot w[m] \cdot e^{-\frac{2\pi j}{M} l(m-k)} \quad (16)$$

where $\mathcal{F}^{-1}\{\cdot\}$ denotes the IFFT operation. Once again, the change of the channel variable from l to k emphasizes the domain transformation. Using the same filter function as before (i.e., $H[l] = A \cdot \delta[l - l_{\text{ch}}]$), the output for each channel is determined:

$$r[n, k] = \frac{A}{M} \sum_{m=0}^{M-1} x[n-m] \cdot w[m] \cdot e^{-\frac{2\pi j}{M} l_{\text{ch}}(m-k)} \quad (17)$$

$$= x[n] * \left(\frac{A}{M} \cdot w[n] \cdot e^{-\frac{2\pi j}{M} l_{\text{ch}}(n-k)} \right) \quad (18)$$

$$= x[n] * h_{\text{ch}}^{(r)}[n, l_{\text{ch}}, k] \quad (19)$$

When $k = 0$, then Equations (12) and (18) are scaled versions of each other by factor M/S_{chan} :

$$y_{\text{chan}}[n] = r[n, 0] \times \frac{M}{S_{\text{chan}}} \quad (20)$$

It shows that the same output result can be achieved but by adding significantly more complexity. All the other output channels (i.e., $k \neq 0$) are phase-delayed replicas.

In the HDDM, the channels are time-delayed and passed through a secondary window $w_s[k]$:

$$r_w[n, k] = r[n - (M - 1) + k, k] \cdot w_s[k] \quad (21)$$

The time delay should counteract the delay introduced for each channel in Equation (1). A fixed delay of $M - 1$ samples is removed to simplify future expressions:

$$\hat{r}_w[n, k] = r_w[n + M - 1, k] = r[n + k, k] \cdot w_s[k] \quad (22)$$

Finally, the outputs are combined:

$$y_{\text{HDDM}}[n] = \frac{1}{S_{\text{hddm}}} \sum_{k=0}^{M-1} \hat{r}_w[n, k] \quad (23)$$

Once again, a constant gain S_{hddm} compensates the processing to achieve unitary gain:

$$S_{\text{hddm}} = \sum_{m=0}^{M-1} w[m] \cdot \sum_{m=0}^{M-1} w_s[m] \quad (24)$$

The output can be determined substituting with Equation (19):

$$y_{\text{HDDM}}[n] = \frac{1}{S_{\text{hddm}}} \sum_{k=0}^{M-1} r[n + k, k] \cdot w_s[k] \quad (25)$$

$$= \frac{1}{S_{\text{hddm}}} \sum_{k=0}^{M-1} (x[n + k] * \left(\frac{A}{M} \cdot w[n + k] \cdot e^{-\frac{2\pi j}{M} l_{\text{ch}} n} \right)) \cdot w_s[k] \quad (26)$$

$$= x[n] * \left(\frac{A}{MS_{\text{hddm}}} \cdot w[n] \cdot e^{-\frac{2\pi j}{M} l_{\text{ch}} n} \right) * w_s[n] \quad (27)$$

$$= x[n] * h_{\text{ch}}^{(\text{HDDM})}[n, l_{\text{ch}}] \quad (28)$$

As with the previous examples, the architecture is band-limited through a filter function $h_{\text{ch}}^{(\text{HDDM})}[n, l_{\text{ch}}]$. The difference between this filter to the previous one is that the window used consists of $2M - 1$ samples through the convolution of the two windows. The approximate doubling in window lengths allows superior side-lobe suppression of the window function. More analysis of the window functions is performed in Section 2.2.

Although only a simple processing function $H_{\text{comp}}(\cdot)$ is used in the examples provided, it demonstrates the single-channel transfer function. The processing function can be significantly more complex. For example, in the interference mitigation case, simple pulse blanking (PB) is commonly used:

$$X_{\text{F}}[n, l] = \begin{cases} X[n, l] & \text{if } |X[n, l]| \leq \lambda_{\text{Th}} \\ 0 & \text{otherwise} \end{cases} \quad (29)$$

where λ_{Th} is the pulse detection and mitigation threshold.

2.1.2 | Multi-rate and batch-processing architectures

Multi-rate processing methods use domain transformations to project a block of data into a different signal

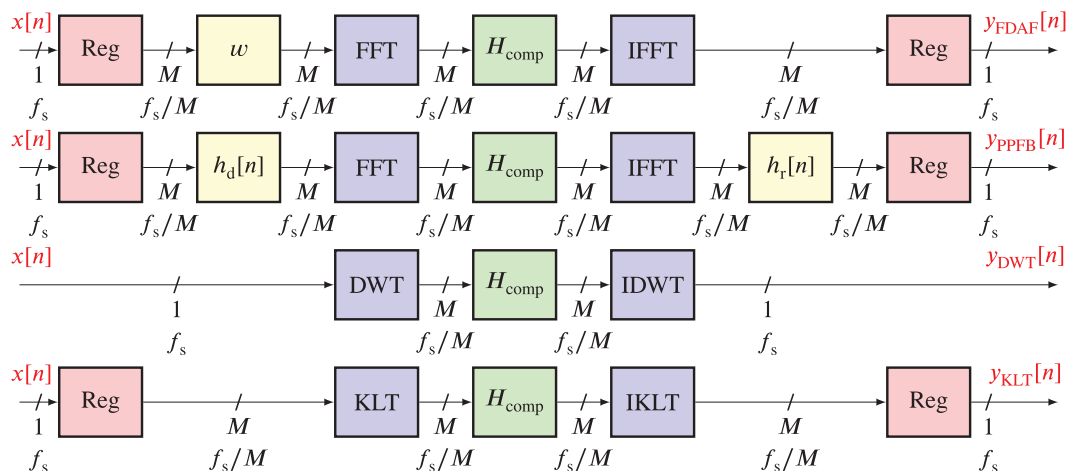


FIGURE 3 Block diagram for multi-rate methods [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

space. The HDDM should be compared to other methods to understand its benefits and shortcomings.

As the processing is implemented on a block of data, an implicit decimation proportional to the block size follows. Some examples include frequency-domain adaptive filtering (FDAF), poly-phase filter bank (PPFB), Karhunen-Loève Transform (KLT), and Discrete Wavelet Transform (DWT) methods (Abdoush et al., 2017; Amin et al., 2017; Borio & Closas, 2019; Burchfield et al., 2020; Chien & Chen, 2015; DAVIS, 2015; Liu et al., 2010; Musumeci & DAVIS, 2012; Querol et al., 2017; Wang et al., 2014). Figure 3 shows the block diagram for these methods. As long as the transformation consists of orthogonal projections, no information is lost. It is the reason why many of these methods are used for audio, image, and video compression.

From an interference mitigation point-of-view, the ideal transformation projects the signals into a signal space where interference is as sparse as possible. Removing the interference will then result in a minimal and insignificant loss to the rest of the signal. However, the reconstruction process of projecting the signal back into the original space can cause deformation issues in signal processing. These are especially problematic when the stochastic behavior of an interference is taken into account. From a restructuring and equalization point-of-view, the transformations should be reversible and accurately represent the imposed deformations.

FDAF is a popular choice for GNSS interference mitigation (Musumeci & DAVIS, 2012; Ojeda et al., 2013; Raimondi et al., 2006). It is similar to the HDDM with no overlap (i.e., each sample is only represented in a single FFT). The FFT size and the alignment of the FFT processing block to the interference have a significant impact on the performance (Raimondi et al., 2006).

For example, consider a band-limited pulse interference with the same duration as the FFT processing block. In the best case, the interference and pulse are aligned, resulting in good characterization by the FFT and consequently interference removal. However, if there is a 50% overlap, the pulse is divided over two FFT blocks which may result in twice as much un-interfered signal being blanked. Alternatively, the interference may not be removed as the spread over two blocks may result in a 3 dB loss of power for detection. This alignment issue is common for all block-wise processing methods. One approach to solve the alignment is to trigger the FFT on a pulsed interference (Gao, 2007). As the HDDM calculates an FFT for each sample offset, this alignment issue is circumvented.

Dual-channel frequency-domain adaptive filtering (Dual-FDAF) (van der Merwe et al., 2018; Zhang et al., 2013) improves the performance by overlapping two FFT blocks. It enhances the temporal properties of the mitigation by suppressing ringing and discontinuities. The HDDM takes this concept to the extreme by maximizing the overlap of each FFT. The performance benefits in interference mitigation also show it (Garzia et al., 2020b).

PPFB, also referred to as *subband adaptive filtering* (SAF) (Borio, 2018; Musumeci & DAVIS, 2012), implements filters with improved properties using more filter coefficients than frequency channels. It improves the temporal scope relative to the FDAF, similar to how the HDDM improves the channelizer.

The PPFB is considerably more flexible than any other mentioned technique, as any arbitrary filter lengths can be realized. Furthermore, through down-sampling before filtering, this approach is highly efficient. The disadvantage is the limited temporal overlap, as is observed with FDAF. To some extent, the HDDM can be considered an

over-sampled filter bank (Wolovich, 1974; Zhou & Do, 2005) with a delay-line reconstruction filter.

The DWT is similar to a DFT, but it uses a different base function called a *wavelet*. A wavelet function facilitates a time-frequency trade-off, allowing more flexibility in signal processing (Daubechies, 1989; Senmoto & Childers, 1972). It is efficient to find structures in a signal and summarize them. Therefore, it is popular for detection, recognition, signal compression, and adaptive resolution media.

Wavelet packet decomposition (WPD) is a widespread and specific implementation. It uses a branch structure, resamples the signals, and finite impulse response (FIR) filters for deconstruction and reconstruction. Many wavelet families exist to choose from, each exhibiting different properties (Mallat, 1989).

The Inverse Discrete Wavelet Transform (IDWT) allows perfect reconstruction of the original signal for several wavelet families. Some GNSS interference mitigation attempts with wavelets have been made (Anyaegbu et al., 2008; Burchfield et al., 2020; Paonni et al., 2010; van der Merwe et al., 2019), but the implementation complexity limits widespread use (Dovis, 2015). However, the DWT is a good analysis and characterization tool.

The KLT transforms the signal to a tailored orthonormal vector set (Szumski, 2011). Therefore, the KLT adapts itself to the signal. Theoretically, the KLT presents the optimal method of decomposing a signal. However, constructing the orthonormal set from the data is highly processing-intensive because it is based on eigenvalue decomposition. The processing requirements make this method impractical (Dovis, 2015).

A method to avoid this is to use *a priori* information to construct the orthonormal set beforehand. However, this approach loses the adaptive properties of the KLT and depends on many assumptions. Further, the KLT shows outstanding detection, characterization, and mitigation performance, but has limited applications beyond the original scope (e.g., it cannot equalize).

Fourier-based multi-rate and batch-processing architectures (i.e., FDAF and PPFb) improve processing efficiency compared to the HDDM. Other transforms (i.e., DWT and KLT) show outstanding performance for detection, characterization, and mitigation, but are still too processing-intensive and have no other signal conditioning benefits. However, the fundamental limitation of all multi-rate methods is the temporal limitation introduced by batch-wise processing and down-sampling.

2.1.3 | Adaptive notch filtering

Adaptive notch filtering (ANF) is a popular architecture used for frequency-modulated continuous-wave (FMCW)

interference signals (Borio, 2016; Borio et al., 2014). It creates a notch over a frequency-sparse signal, but it can also change the notch as the signal changes. The infinite impulse response (IIR) implementation of the ANF is particularly attractive as it requires low resources. A limitation is that this approach is only usable against these signals and results in degradation for other interference signals (e.g., wideband noise signals or multiple simultaneous signals) (van der Merwe et al., 2020).

Furthermore, a limitation of ANF is that it is only applicable to interference mitigation and not other signal conditioning functions. A recent study by Borio et al. showed that ANF causes significant biases in the pseudo-ranges (Borio & Gioia, 2020), which degrade the position, velocity, and time (PVT) solution. Therefore, ANF is a poor signal conditioning approach despite its good interference suppression capabilities for select interferences.

2.2 | Theoretical analysis of the HDDM

The band-pass filters of Equations (13) and (28) each consist of three parts: First, the gain A determines the through-gain of the channel; second, a complex exponential $e^{-\frac{2\pi j}{M}l_{ch}n}$ selects the center frequency of the filter band; and third, a window function determines the pass-band bandwidth and the side-lobe suppression. When comparing the filters directly to each other, it is clear that the only difference is the window function:

$$h_{ch}^{(\text{chan})}[n, l_{ch}] = \frac{A}{S_{\text{chan}}} \cdot e^{-\frac{2\pi j}{M}l_{ch}n} \cdot w[n] \quad (30)$$

$$h_{ch}^{(\text{HDDM})}[n, l_{ch}] = \frac{A}{MS_{\text{hddm}}} \cdot e^{-\frac{2\pi j}{M}l_{ch}n} \cdot (w[n] * w_s[n]) \quad (31)$$

The selection of the window functions $w[n]$ and $w_s[n]$ determine the performance of the two algorithms. In both cases, the first window determines how the processing function $H_{\text{comp}}(\cdot)$ responds to the signal.

For example, in the PB case, only this window influences whether a channel is blanked. Therefore, the first window is crucial for detection, classification, estimation, and initial processing. The second window $w_s[n]$ of the HDDM determines the signal reconstruction. It affects the entire system response, but it does not affect any processing methods before the IFFT.

There are many window functions (Ifeachor & Jervis, 2002), and it is thoroughly researched in the literature, but for simplicity, only rectangular (boxcar) and Hann windows are considered for this discussion. Window functions provide a trade-off between the width of the main-lobe, suppression of side-lobes, and the location of zeros in the spectrum. A rectangular window has the narrowest, ideal




FIGURE 4 Window function over time [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]




FIGURE 5 Frequency response of window functions [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

main-lobe properties and well-placed zeros (the rectangular window results in a sinc function in the frequency domain where zeros ensures ideal frequency-bin orthogonality), but it has inferior side-lobe suppression. A Hann window has an increased main-lobe width and non-ideal zero locations but has superior side-lobe suppression.

Figure 4 shows the time response of some window functions, while Figure 5 shows their spectrum responses.

For example, at 2.5 bins from the center, a rectangular window suppresses 8.8 dB and the Hann-32 window 15.8 dB. Both the Hann-64 and HDDM-32 windows (consisting of the joint rectangular and Hann-32) have a suppression better than 24 dB. It shows that the HDDM-32 has similar suppression capabilities to the Hann of the same length. However, as the frequency bins increase,




FIGURE 6 Combined frequency response of the entire HDDM [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

the HDDM shows superior suppression capabilities. The HDDM-32 has the same number of channels as the Hann-32 but has, at 2.5 bins, a 10 dB amplitude isolation improvement. Furthermore, the HDDM has a main-lobe width similar to the rectangular window. It means that the increased window length of the HDDM implicitly improves both the side-lobe suppression and main-lobe width compared to a channelizer architecture with the same FFT size.

Figure 6 shows the response of the windowed channels (solid colored lines) and the combined response (black dashed line). It shows the theoretical pass-band ripple to be below 0.01 dB, which is negligible.

2.3 | Complexity analysis

The number of multiplications $N_{\text{mult}}^{(\text{HDDM})}$ and additions $N_{\text{add}}^{(\text{HDDM})}$ per clock-cycle for the HDDM relate to the FFT size M (assuming it is based on the efficient Radix-2 implementation):

$$N_{\text{mult}}^{(\text{HDDM})} = 2M \log_2(M) \quad (32)$$

$$N_{\text{add}}^{(\text{HDDM})} = M + 2M \log_2(M) \quad (33)$$

Similarly, the number of multiplications $N_{\text{mult}}^{(\text{chan})}$ and additions $N_{\text{add}}^{(\text{chan})}$ for a channelizer (i.e., not using the IFFT stage to reconstruct the signal) is also determined:

$$N_{\text{mult}}^{(\text{chan})} = M \log_2(M) \quad (34)$$

$$N_{\text{add}}^{(\text{chan})} = M + M \log_2(M) \quad (35)$$

TABLE 1 Complexity comparison

Architecture	DFT	Filter	Operations	
	M	K	N_{mult}	N_{add}
Chan-32	32	–	160	192
HDDM-32	32	–	320	352
Chan-64	64	–	384	448
FDAF-32	32	–	10	10
PPFB-32	32	64	14	12

These expressions do not include the multiplications of the window functions, as these could add zero complexity for rectangular windows. Simple bit-shifts can also be used for windowing to avoid multiplications (Garzia et al., 2020a). If a window function is multiplied, then additional M multiplications are required. The complexity of the processing function $H_{\text{comp}}(\cdot)$ is also omitted, as this would change depending on the application. For example, simple PB requires only basic logic functions and no additions or multiplications. In another example, equalization (see Section 6) requires an additional M complex multiplication.

The multi-rate methods operate batch-wise, which requires significant memory buffering. With every M clock-cycle a new batch is processed. The FDAF algorithm requires factor M fewer operations than the HDDM (including windowing, but excluding the reconstruction additions after the IFFT):

$$N_{\text{mult}}^{(\text{FDAF})} = 2 \log_2(M) \quad (36)$$

$$N_{\text{add}}^{(\text{FDAF})} = 2 \log_2(M) \quad (37)$$

A PPFB is more complex than FDAF due to the added filter. The PPFB has an FFT of size M and a filter consisting of K coefficients:

$$N_{\text{mult}}^{(\text{PPFB})} = 2 \frac{K}{M} + 2 \log_2(M) \quad (38)$$

$$N_{\text{add}}^{(\text{PPFB})} = 2 \frac{K - M}{M} + 2 \log_2(M) \quad (39)$$

As seen in both the FDAF and the PPFB, these require significantly fewer processing resources. However, these two methods are limited by the increased memory resources, limited parallelization due to the batch-wise processing, reduced temporal resolution, and are less versatile. The DWT and KLT are omitted due to the increased complexity of the design choices. However, both require significantly more resources than the FDAF.

Table 1 compares the complexity of the HDDM, channelizer, FDAF, and PPFB. A channelizer with the same number of channels uses approximately half of the resources compared to the HDDM. This difference is due to the additional IFFT stage. However, the HDDM has superior isola-

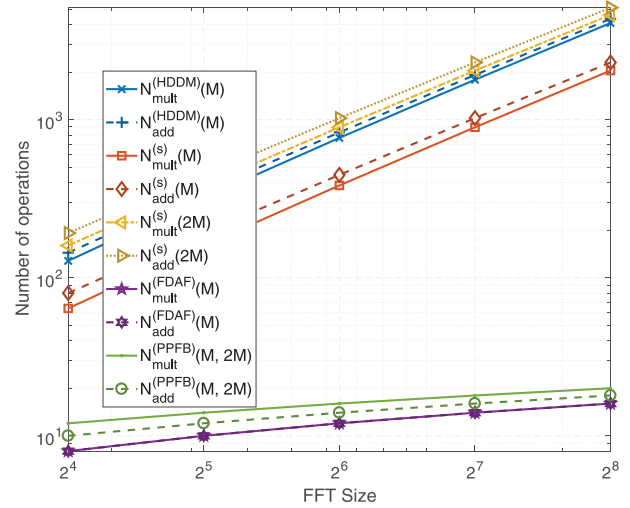


FIGURE 7 Number of operations [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

tion capabilities due to the double window function, as is shown in the previous section. Therefore, it has similar isolation capabilities to a channelizer with double the amount of channels. The HDDM is less complicated than a channelizer with double the number of channels and consequently is a trade-off between the complexity and isolation capabilities compared to a channelizer architecture. The spectrum isolation capabilities are more important than the complexity for interference mitigation, resulting in the HDDM architecture niche.

By far, the FDAF is the least resource-intensive method (excluding memory restrictions), using more than an order of magnitude fewer resources than the HDDM. However, there is a significant performance difference between the FDAF and HDDM (van der Merwe et al., 2020a). Most notably is the temporal resolution of the HDDM, which gives it superior performance to pulsed interferences. The PPFB uses slightly more resources than the FDAF but has superior side-lobe suppression. It is a highly efficient method, which clearly emphasizes its suitability for compression applications. However, it has the same time limitations as the FDAF.

Figure 7 shows the number of operations for different designs. There is a great complexity difference between the multi-rate methods and the HDDM. However, the versatility of the HDDM is regarded in this section. If the HDDM signal conditioning fulfills multiple roles, it may be more efficient than a discrete processing chain.

For example, using a multi-rate technique only for interference mitigation, coupled with a purposefully developed FIR filter for equalization, and several FIR filters and mixers for spectrum restructuring. Another significant aspect not discussed is the memory requirements. Batch

processing methods are efficient in software processing [i.e., software-defined radio (SDR)], as large memory resources are available on most computer systems. However, in optimized firmware and hardware systems for low SWAP-C embedded applications, memory is costly. In the worst case, logic space is purposed for memory on field-programmable gate array (FPGA) systems. Hence, the target system and the memory availability have a significant impact on algorithm selection.

3 | HARDWARE IMPLEMENTATION

The first hardware implementation realization of the HDDM was used for interference mitigation (Garzia et al., 2020a). In this case, an HDDM-32 equipped with PB was implemented on a wideband (WB) high-end GNSS receiver. Afterward, the HDDM-32 was ported to low-end FPGA-based GNSS receivers. A different mitigation approach has also been evaluated, where the PB is replaced with automatic gain control (AGC) modules (Garzia et al., 2021). Therefore, maintaining the GNSS signals in tracking for continuous interferences like continuous-wave (CW) or WB noise.

The HDDM hardware module has been designed in Register-Transfer Level (RTL) code [i.e., very-high-speed integrated circuit hardware description language (VHDL)], allowing the implementation on different FPGA technologies or even as an application-specific integrated circuit (ASIC) intellectual property (IP).

According to Figure 2, the HDDM can be divided into three parts: data deconstruction, data manipulation, and data reconstruction. The deconstruction is composed of a sequence of delay stages (same amount as FFT points, i.e., 32) which are realized using the flip-flops (FFs) in the FPGA fabric.

In this stage, for each input sample, we obtained 32 output samples. The 32 samples were processed in parallel in the following stages. A standard rectangular window characterizes the windowing stage (in yellow in the block diagram). However, instead of multiplying the samples by a set of coefficients, we performed some combination of bit shifting. This way, resource-expensive multipliers are replaced by simple bit manipulations, avoiding sparse digital signal processor (DSP) blocks in FPGA implementations.

The final stage of the deconstruction was a 32-point FFT. The architecture of the FFT was based on a Radix-2 Cooley-Tukey algorithm. The five stages of the FFT were completely deployed in the FPGA fabric: trading flexibility for throughput and processing 32 samples in one clock cycle per stage. Pipelining between the FFT stages facilitated the final throughput of 32 samples per cycle.

The FFT core needs as many multipliers as the algorithm requires since no resource sharing is done. However, some optimizations are still possible: the twiddle factors are hard-coded in hardware (HW), which means that the multiplications were between an input and a constant instead of two inputs. It significantly reduces the HW resources and allows for some resource sharing at the DSP block level.

The data manipulation depends on the specific use-case. In the case of PB interference mitigation, it is composed of comparators that check the samples (FFT outputs) against a threshold and set them to zero when above the threshold. The comparators were implemented using look-up tables (LUT). In the case of spectrum restructuring and equalization, DSP blocks might be needed. There are mixer algorithms that do not need any multiplier, and equalization can be done on a power-by-two basis, which only uses bit shifting [i.e., the Coordinate Rotation Digital Computer (CORDIC)].

In the case of AGC-driven mitigation, a resource-efficient bit-shift is implemented with a varying amount of shifts. A control logic circuit calculates the average signal amplitude and selects the most suitable shifted version of the input signal. The benefit is that no multiplication is used, which spares DSP slices on FPGA implementations.

The final reconstruction stage is composed of an IFFT which is implemented in a similar way as the direct one, but optimized for the inverse algorithm to limit resource utilization. The throughput is 32 samples per clock cycle as for the FFT.

The final stage is comprised of a sequence of adders and delay blocks that recombine the 32 IFFT outputs described in Section 2. The structure is reorganized so that there is a maximum of one adder between two registers: this reduces the critical timing path (see Figure 8). The whole structure is pipelined to process 32 samples per clock cycle and produces one sample per cycle, equivalent to the HDDM input.

The whole structure can be scaled up to a larger amount of FFT points to increase frequency selectivity. However, due to the optimized structure, the FFT core would need a design adaptation. The resource utilization would then also scale accordingly.

The implementation details including bit widths and maximum operating frequencies vary between use-cases and target HW platforms. One aspect that plays an important role is the input bit width [which is related to the analog-to-digital converter (ADC) used in the platform] and has an impact on the interference mitigation capabilities. An evaluation of the interference mitigation performance for different bit widths is provided in Garzia et al. (2020b).

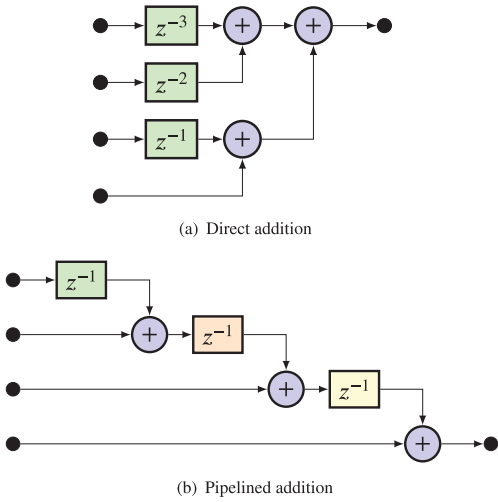


FIGURE 8 Comparison of triangular register methods [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]


4 | INTERFERENCE MITIGATION

The original purpose of the HDDM algorithm was interference mitigation (van der Merwe et al., 2020a). Interference mitigation aims to remove interference signals such that the receiver can optimally process the GNSS signals.


4.1 | Interference mitigation case study

Any non-GNSS signal in the GNSS spectrum is considered an interference, as it would degrade GNSS performance. However, different interferences have different mitigation efficiency based on the mitigation process.


Consider the four signals in Figure 9 to illustrate the mitigation process. A WB receiver with 40 MHz sample rate is simulated. PB and ANF are two extreme cases for interference mitigation, and the HDDM—or any other multi-spectral method—is a trade-off between the two.




(a) Single tone CW



(b) Chirp signal



(c) Pulsed wideband noise



(d) Frequency hopper

FIGURE 9 Interference scenarios [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

Figure 9(a) is a simple CW signal. ANF—or even just regular notch filtering—will mitigate this signal without a problem. However, PB will either do nothing or remove the entire signal due to the constant nature of the signal. As a result, it can only degrade the signal. Any multi-spectral method (e.g., HDDM or FDAF) will also be able to remove this signal. However, the frequency resolution will be a limiting factor for most cases.

Figure 9(b) is a chirp signal where the signal changes over time. PB will have the same issues as in the previous case. A notch filter will no longer suffice, and ANF will be necessary to adapt to the signal. If the ANF is correctly configured, then the only issue would be the jump discontinuity when the chirp jumps from the higher band-edge to the lower one.

The multi-spectral techniques will also be able to mitigate this interference. However, the block size or FFT length directly influences the performance. For example, consider the FDAF with an FFT exceeding 20 μ s of data, then an entire cycle of the chirp is enclosed in one FFT and runs the risk of removing the entire signal. It will degrade the performance significantly.

If the FFT is too small, then the spectral resolution is also lost and more bandwidth per FFT is removed than necessary. The HDDM facilitates a better time-frequency trade-off due to the oversampling of the FFT, which will contribute to improved performance.

Figure 9(c) shows a WB pulsed signal. PB will be able to mitigate this interference successfully, resulting in an ideal degradation of 3 dB, as half of the signal power is removed along with the noise interference. ANF will not work in this case, as the signal is not sparse in the frequency domain.

It illustrates the trade-off between PB and ANF based on the interference signal type. Multi-spectral techniques would remove this interference, but once again, it is granular to the FFT size and the block offset. For example, if the processing block is the same size as the pulse and aligned with the pulse, then FDAF could present the same performance as PB. However, if the block is delayed with half a pulse, then each block contains half, resulting in either removing the entire signal or no mitigation. With shorter block sizes, the time granularity may be improved, but at the cost of frequency resolution.

The HDDM does an FFT for each offset, which means that it is guaranteed to be aligned with the pulse edge at a given time, resulting in superior performance compared to the FDAF. However, the successive FFTs in the HDDM result in a low transition, so the pulse edge is no longer pronounced.

Finally, consider Figure 9(d). If this interference were in the L1 frequency band, the pulses would interfere with the upper and lower side-bands of higher-order binary offset carrier (BOC) signals such as the Global Positioning Sys-

tem's (GPS) M-Code or Galileo's E1A. PB will not be possible in this case, whether it removes the entire signal or does nothing, as the signal is continuously transmitting. This interference shows an example where both PB and ANF will be ineffective. FDAF will have the same limitations as in the previous cases and the HDDM will have the same advantages.

4.2 | Interference mitigation performance

An excerpt of previous research results is presented (Garzia et al., 2020b) to demonstrate the interference mitigation performance. Four interference signals are selected:

1. Single tone CW at 1,575.42 MHz
2. Two equal-amplitude simultaneous chirps, both with 35 MHz bandwidth, but with 16 μ s and 100 μ s repetition rates, respectively
3. A pulsed noise signal, with 35 MHz bandwidth, 0.5 ms pulse width, and 50 % duty cycle
4. A frequency hopper with a dwell time of 1 μ s and a frequency range of 35 MHz

The receiver-estimated carrier-to-noise density ratio (C/N_0) for Galileo's E1BC from three GNSS receivers are compared:

1. A WB GNSS receiver with an HDDM-32 and PB implemented in firmware (see Section 3 for the implementation details) with a sample rate of 108 MHz with a 12-bit ADC.
2. A mass-market commercial-off-the-shelf (COTS) receiver
3. A high-end COTS GNSS receiver with known good interference mitigation capabilities

The interferences start with low power, increasing with 1 dB every 10 s over a range of 70 dB. At this stage, most receivers cannot track the satellite signals. Then the power reduces at the same rate back to the starting value. The receivers have different tracking and acquisition sensitivities and will lose lock at different C/N_0 values. However, the general trends are visible.

Figure 10 shows the C/N_0 and the power of the four interferences. The HDDM has marginally improved performance against the double chirp (Figure 10b) and frequency hoppers (Figure 10d) compared to the current high-end state-of-the-art. Further, it excels against pulsed interferences (Figure 10c) and outperforms all other methods. However, it suffers from single-tone CW interferences (Figure 10a) due to the limited frequency resolution.

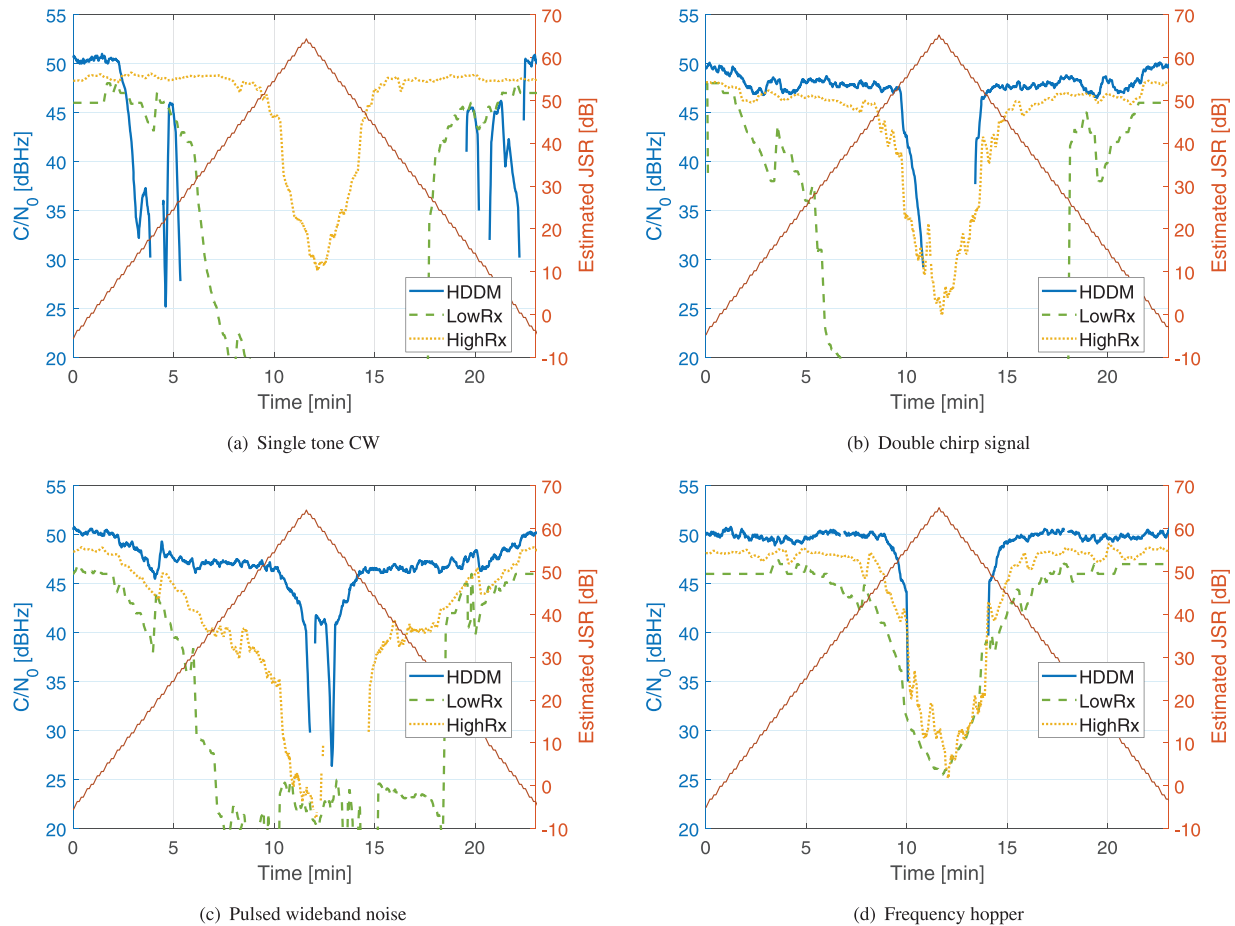


FIGURE 10 Interference mitigation performance comparison, based on Garzia et al. (2020b) [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

A method to address this limitation is to combine the HDDM with a simple notch filter. The HDDM and notch filter combination proved to be a good synergy in previous studies (Garzia et al., 2020b).

5 | SPECTRUM RESTRUCTURING

The FFT in the deconstruction stage channelizes the input signal. One option is to filter and frequency-translate the GNSS signals to alter their spectral properties (van der Merwe et al., 2020b). Restructuring results in two major benefits: Compressing the spectrum allows for lower sample rates and, consequently, lower data throughput. It results in a significant reduction in storage and processing requirements for SDR receivers. Reduced processing requirements also facilitate low SWAP-C for conventional receivers. It also facilitates joining different frequency bands together.

If the HDDM is already implemented for interference mitigation, there are synergies when combining with spectrum reconstruction. The high dynamic range required for

interference mitigation is no longer needed after mitigation. Hence, both the sample rate and the word-length are reduced simultaneously. This reduction has significant benefits for SDR applications. The second synergy is that the same hardware resources are used. Hence, the additional resources needed for spectrum reconstruction are insignificantly small. This aspect emphasizes the main benefit of the HDDM architecture: Multiple signal conditioning methods simultaneously reduce the overall resource requirements.

Figure 11 shows an example of the L1 frequency band in which the spectrum is filtered and restructured using the HDDM. The original signal is sampled at 81 MHz with 8 bits containing all L1 open service signals: GPS's L1 C/A, Galileo's E1B/C, Beidou's B1I, and GLONASS's G1. An HDDM with a 32-point FFT is used, which relates to a spectral granularity of 2.53 MHz. The final output spectrum preserves the signals and compacts these to occupy less than 40 MHz. The measured PSD (the thick blue lines at the top of Figure 11) and the band-limited theoretical PSDs (the colored thin lines beneath) are shown. The reconstruction is done by reassigning the frequency channels




FIGURE 11 Spectrum restructuring example [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

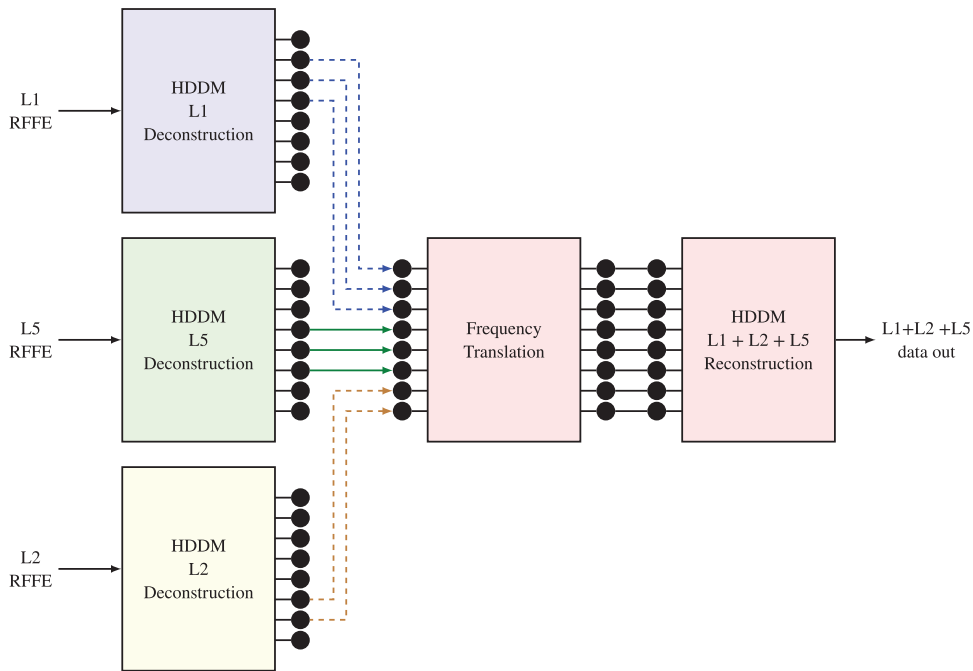


FIGURE 12 Block diagram for an overlay design [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

(i.e., p translates to l), and adding a mixer to each channel:

$$X_F[n, l] = H_{\text{comp}}(X[n, p]) = e^{-2\pi j n \times \frac{l-p}{M}} \times X[n, p] \quad (40)$$

The HDDM already implicitly filters the signal into multiple channels, so no further filtering is required. Therefore, spectrum reconstruction requires few additional resources.

Alternatively, multiple frequency channels can be joined together (Rügamer et al., 2020). Figure 12 demonstrates such a design and shows the manipulated spectrum and the theoretical PSDs.

Three frequency bands are sampled and passed through HDDM deconstruction stages. Next, the channels are interleaved, frequency translated, and reconstructed using a single HDDM reconstruction stage. It results in a single,

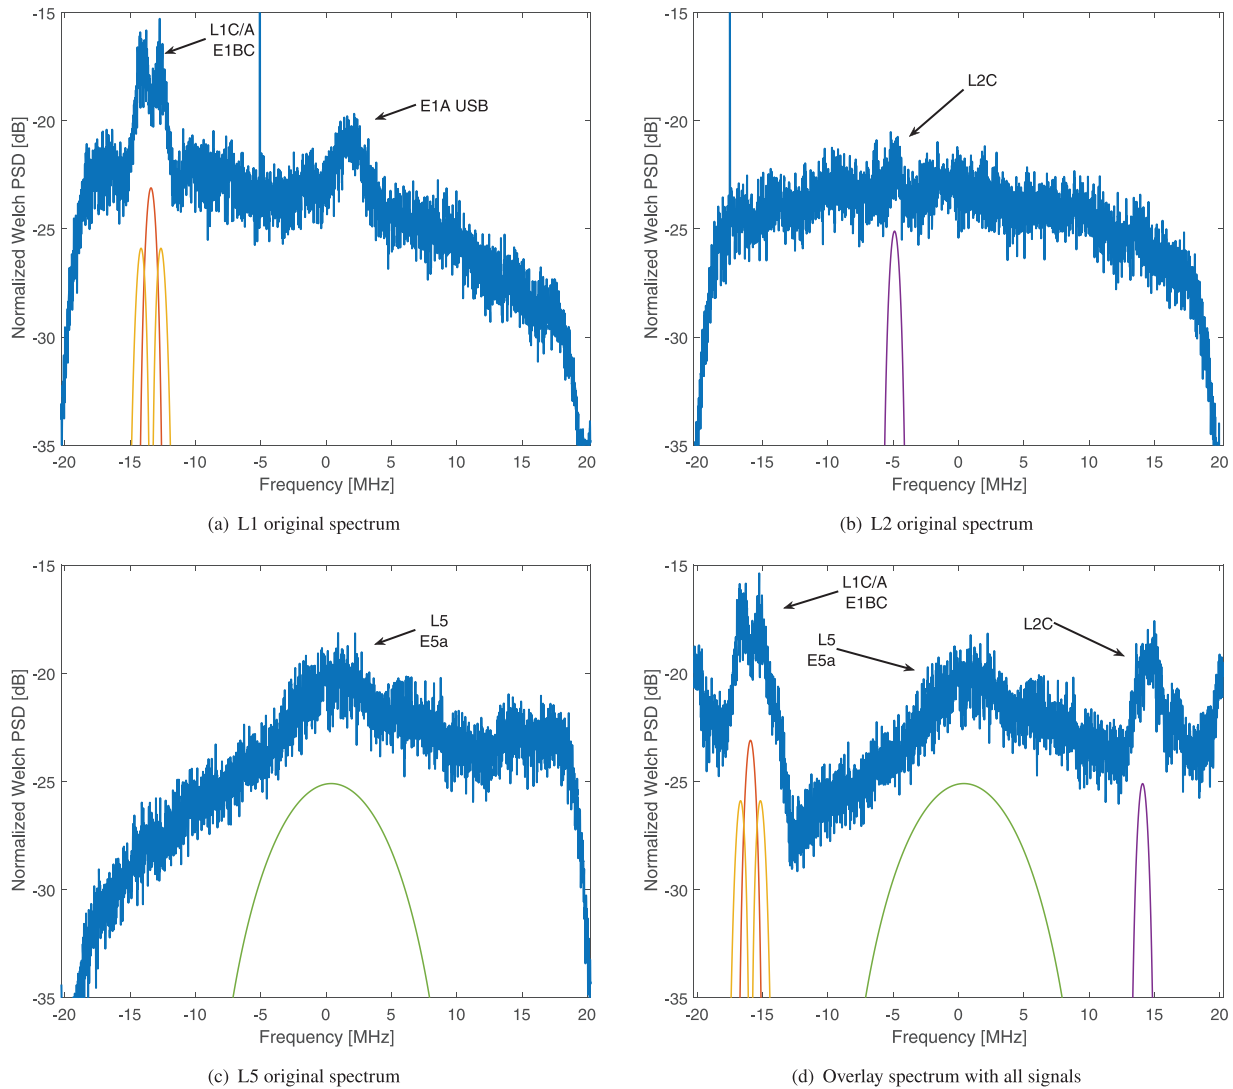


FIGURE 13 Power spectral densities (PSDs) for overlay spectrum combinations [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

compact data stream containing signals from multiple frequency bands.

Figure 13 shows the PSDs for a practical example of the overlay design. A triple-band receiver records L1, L2, and L5 signals at a sample rate of 40.5 MHz with 4-bit IQ. The spectra are filtered, frequencies translated, and combined into a single data stream. An HDDM with a 32-point FFT is used, which relates to a spectral granularity of 1.27 MHz.

The resultant data allows triple-band signal tracking to facilitate multi-band processing within a single, artificially reconstructed, compact band. Furthermore, the data rate is reduced by a factor of three, resulting in significant storage and transmission reductions. Lastly, this overlay approach is an interesting pre-processing stage to facilitate efficient meta-signal processing (García-Molina & Fernández-Rubio, 2019; Paonni et al., 2014).

The spectra are combined digitally, which means that the typical anti-aliasing filter roll-off is not needed. It is shown in Figure 13(d) where the entire spectrum is fully used. Furthermore, aliased signals, which cross over the Nyquist rate, are not an issue. These signals can be easily shifted to a zero intermediate-frequency (IF) before GNSS processing. Therefore, an overlay design results in efficient spectrum utilization.

The overlapping spectra cause self-interference, and consequently, a loss in the received C/N_0 (Rügamer et al., 2020). However, through the HDDM-induced filtering, interference is suppressed. An optimized overlay design can adequately reduce the C/N_0 loss to only one or two decibels. Furthermore, the code-division multiple access (CDMA) signal structure used by GNSSs provides some jamming-resistant qualities. Therefore, despite these degradations, good signal tracking is expected, and the

benefits of reduced data rates and multiple frequency band operation outweigh these limitations.

6 | EQUALIZATION

The satellite transmission equipment, link between the satellite and receiver, and receiver hardware degrade the received signal (Rappaport, 1996). This degradation results in changes in the signal magnitude and phase for different spectral components. Signal equalization can correct these deformations and counter the channel effects. Equalization is commonplace in communication systems (Lathi & Ding, 2010). However, due to the low received signal power, the CDMA structure used for GNSS, and the lack of sounding signals, it is impractical to equalize the transmission channel for each satellite signal. On the contrary, the benefit of the CDMA structure is the additional robustness against multipath and fading effects. Further, great care is taken to ensure that the transmission hardware is of high quality. Only the receiver hardware remains a dominant source for signal degradation.

Low-cost mass-market receivers usually have poor front-end characteristics, including poor antenna design, which greatly deform the received signals. However, medium- to high-end receivers are also influenced by the front-end designs (Guerrero & Gunawardena, 2017). Surface acoustic wave (SAW) filters are commonly used for GNSS receivers due to their small size and excellent frequency selectivity. The selectivity is crucial to reduce interference and inter-modulation products from other frequency bands, as other signals are typically received at much greater power than GNSS signals.

The disadvantage of SAW filters is that they introduce significant phase distortions and group delays to the received signals. It has been shown problematic for higher-order BOC modulated signals (Liu et al., 2020; Wendel et al., 2019). These group delays result in problems during tracking, introduce pseudo-range biases, and degrade the quality of the PVT. If the front-end is characterized, then the phase and magnitude response can be equalized using the HDDM.

The limited frequency resolution of the HDDM restricts the equalization capabilities. Further, the side-lobes introduced by the window function also result in some distortion in the frequency domain.

As an example of the equalization capabilities, an HDDM-32 at a sample rate of 81 MHz is simulated. An RFFE containing a Vectron TFS1575T SAW filter is assumed. This filter has a center frequency of 1,575 MHz and a specified 1 dB analog bandwidth of 30 MHz or 50 MHz 3 dB bandwidth. Two equalization filters are generated by adapting the complex gain of each HDDM




FIGURE 14 Magnitude response of the front-end filter [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

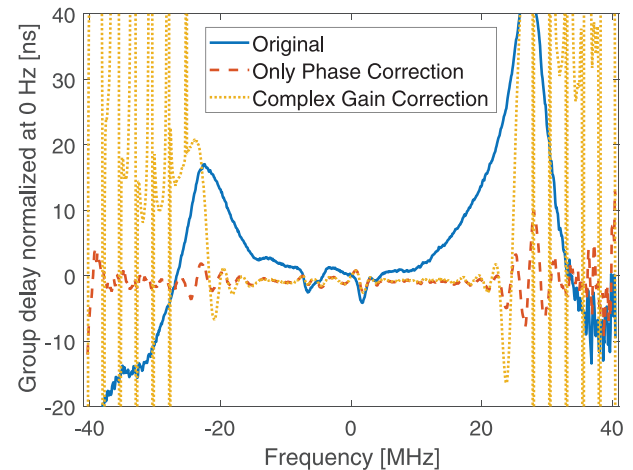


FIGURE 15 Group delay of the front-end filter [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

channel. The first filter only corrects the phase offsets. The second filter corrects the phase and magnitude. The phase-only design uses a low-complexity CORDIC to correct the phase to minimize multiplications. On the contrary, the amplitude and phase design requires a complex multiplication with significantly more hardware resources.

Figure 14 shows the gain of the SAW filter and the corrected gains after equalization. The SAW filter has a flat response. Hence, limited gain corrections are necessary. However, out-of-band frequencies are further suppressed by the HDDM to limit interference. Figure 15 shows the group delay. The SAW filter has a group delay difference of up to 20 ns in the pass-band. After equalization, it reduces below 5 ns. A group delay of 20 ns is not a primary

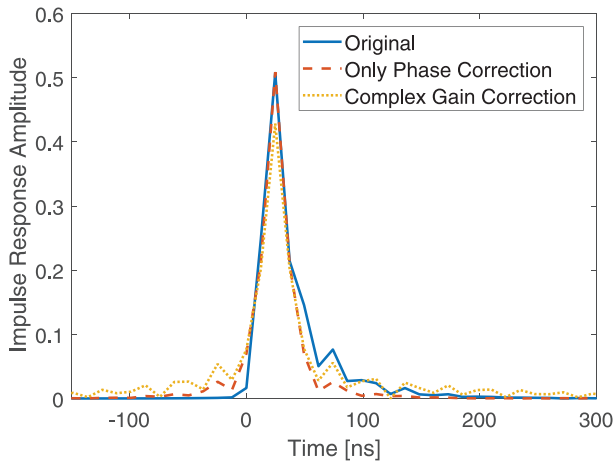


FIGURE 16 Impulse response of the front-end [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

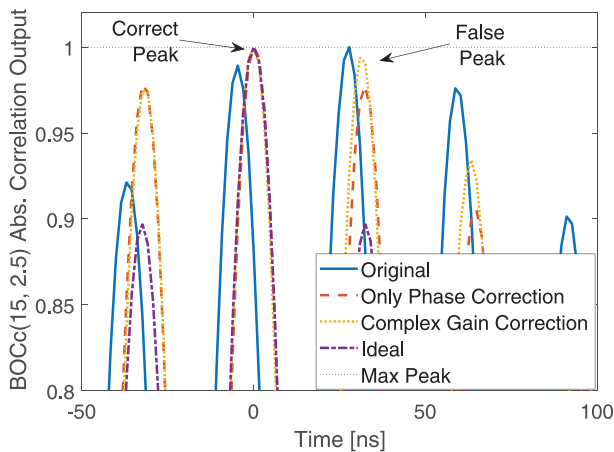


FIGURE 17 BOCc(15, 2.5) ACF distortion of the front-end [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

concern for most GNSS signals, but it can still impact the receiver accuracy.

Figure 16 shows the impulse response, and Figure 17 shows the impact of the filters on the auto-correlation function (ACF) of a BOCc(15,2.5) signal. The impulse response has some amplitude values preceding the main peak and indicates the effort of the equalization to correct the signal. This signal's modulation is commonly used for group delay analysis in GNSS receivers (Liu et al., 2020; Wendel et al., 2019) due to its high bandwidth requirements and challenging multiple correlation peaks to process. The SAW filter causes the first trailing side-peak of the ACF to be larger than the actual one.

As a result, a GNSS receiver would most likely lock onto this peak, resulting in pseudo-range biases that can cause sub-optimal PVT solutions. The two equalization filters




FIGURE 18 Spectrum before and after amplitude equalization [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

suppress this side-peak such that it is not larger than the correct one. However, the filters do not suppress them as low as the ideal side-peak. This deficit shows the limitations of equalization capabilities of the HDDM, but it does correct the group delays to impact receiver accuracy significantly.

Amplitude irregularities degrade the correlation peak and incur tracking instabilities for wideband signals and higher-order BOC signals. For example, a greater than 10 dB amplitude difference between the upper and lower side-band of a BOC signal is approximately the same as single side-band tracking. It causes significant vestigial signals in the quadrature-phase due to the BOC to binary phase-shift keying (BPSK) mismatch.

Figure 18 shows some practical results for amplitude corrections only. The receiver has a sample rate of 108 MHz, a center frequency of 1,575 MHz, and uses an HDDM-32. The amplitude corrections compose of only bit-shifts (i.e., corrections are exponents of two), resulting in a hardware-efficient approach. An AGC processor adapts the amplitude gain for each channel independently. As seen, the spectrum is flattened out, countering the RFFE effects. It is especially evident in Galileo's BOCc(15,2.5) and BeiDou's BOC(14,2) signals, where a 2 dB difference between the lower side-band is corrected (at -15 MHz).

7 | DISCUSSION AND OUTLOOK

The HDDM is a versatile digital signal conditioning architecture. However, the preceding stages in a receiver should also be considered for optimal design (i.e., GIGO). For example, if the RFFE is adequately designed, then

equalization with the HDDM plays a more minor role. Another example is that antenna interference suppression techniques may have superior performance and may save later stages from unrecoverable saturation.

Ideally, the entire processing chain should be considered as a whole. One proposal is to use the RFFE's AGC information to adjust the interference mitigation thresholds inside the HDDM. Such an approach would result in optimal threshold setting and performance. Another proposal is to incorporate a calibration signal in the receiver to characterize the RFFE and equalize the receiver dynamically. However, such holistic design approaches depend on the design philosophy of the receiver development team.

Restructuring the spectrum is a particularly interesting use of the HDDM. Several examples of spectrum compression and overlay designs have been shown in this paper. However, a systematic characterization of such an overlay design is proposed for future work. The potential to have more compact data streams has significant benefits for digital storage and transmission. Therefore, SDR receivers (Pany, 2010) and snapshot-based receivers (Lucas-Sabola et al., 2016; Rubino et al., 2016; Rügamer et al., 2016) could see significant improvements. Furthermore, restructuring the spectrum may also apply to altering a WB signal's properties. For example, a higher-order BOC signal can be re-modulated as a lower-order BOC [e.g., a BOCc(15, 2.5) re-modulated as a BOC(8, 2.5)].

Lastly, the HDDM shows superior time-frequency trade-offs in processing. The concept may be expanded to other domains for future research. By using the HDDM with a multi-antenna receiver, space-frequency-time manipulation could be possible. It is currently theoretical but could be an exciting new avenue of research.

8 | CONCLUSION

The HDDM facilitates advanced multi-functional signal conditioning—including interference mitigation, spectrum restructuring, and equalization. Each of these cases is showcased in this article with theoretical designs and practical examples. It displays the versatility of the method and efficient joint processing. Several applications of the HDDM are explored in this article, and others are proposed for future research. An important aspect is that the HDDM can simultaneously do several signal conditioning processes, saving on overall system processing requirements.


This paper presents the first theoretical modeling of the HDDM for GNSS applications. Theoretical modeling shows the fundamental feature of the method: an increased temporal window range with increased temporal resolution. The HDDM has superior spectral properties

to channelizer and FDAF architectures and superior temporal properties to FDAF, PPFB, and other multi-rate methods. A limitation is that it is more processing-intensive than FDAF or PPFB architectures, but the multi-functionality and increased performance justifies this deficit. Lastly, an example hardware implementation showed that the HDDM can be practically implemented in a GNSS receiver.


CONFLICT OF INTEREST

The authors declare no potential conflict of interests.

ORCID

Johannes Rossouw van der Merwe  <https://orcid.org/0000-0002-8441-5515>

Fabio Garzia  <https://orcid.org/0000-0003-2597-4429>

Alexander Rügamer  <https://orcid.org/0000-0002-5927-4685>

Wolfgang Felber  <https://orcid.org/0000-0003-1927-7870>

REFERENCES

- Abdoush, Y., Pojani, G., Bartolucci, M. & Corazza, G. E. (2017). Time-frequency interference rejection based on the S-transform for GNSS applications. *2017 IEEE International Conference on Communications (ICC)*, 1–6. <https://doi.org/10.1109/ICC.2017.7996554>
- Amin, M. G., Borio, D., Zhang, Y. D., & Galleani, L. (2017). Time-frequency analysis for GNSSs: From interference mitigation to system monitoring. *IEEE Signal Processing Magazine*, 34(5), 85–95. <https://doi.org/10.1109/MSP.2017.2710235>
- Anyaegbu, E., Brodin, G., Cooper, J., Aguado, E., & Boussakta, S. (2008). An integrated pulsed interference mitigation for GNSS receivers. *Journal of Navigation*, 61(2), 239–255. <https://doi.org/10.1017/S0373463307004572>
- Borio, D. (2016). Loop analysis of adaptive notch filters. *IET Signal Processing*, 10(6), 659–669. <https://doi.org/10.1049/iet-spr.2015.0310>
- Borio, D. (2018). Sub-band robust GNSS signal processing for jamming mitigation. *2018 European Navigation Conference (ENC)*. Gothenburg, Sweden, 72–83. <https://doi.org/10.1109/EURONAV.2018.8433250>
- Borio, D., & Closas, P. (2019). Robust transform domain signal processing for GNSS. *NAVIGATION*, 66(2), 305–323. <https://doi.org/10.1002/navi.300>
- Borio, D., & Gioia, C. (2020). Gns interference mitigation: A measurement and position domain assessment. *NAVIGATION*, 68(1), 93–114. Retrieved from <https://doi.org/10.1002/navi.391>
- Borio, D., O'Driscoll, C. & Fortuny, J. (2014). Tracking and mitigating a jamming signal with an adaptive notch filter. *Inside GNSS*, 67–73. <https://www.insidegnss.com/auto/marapr14-WP.pdf>
- Bruton, A. (1997). Reduction of GPS receiver noise using adaptive filters. *Proc. of the 10th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 1997)*. Kansas City, MO, 645–663. <https://www.ion.org/publications/abstract.cfm?articleID=2856>
- Burchfield, S., Martin, S., Bevely, D. & Starling, J. (2020). Performance analysis of low SWaP-C jamming mitigation methods for commercial applications. *Proc. of the 33rd International Technical Meeting*

- of the Satellite Division of The Institute of Navigation (ION GNSS+ 2020), 3592–3611.
- Chien, Y. & Chen, P. (2015). Wavelet-packet-transform-based adaptive predictor to mitigate GNSS jammers. *2015 International Conference on Wavelet Analysis and Pattern Recognition (ICWAPR)*, 14–19. <https://doi.org/10.1109/ICWAPR.2015.7295919>
- Daubechies, I. (1989). Wavelets: a tool for time-frequency analysis. *Sixth Multidimensional Signal Processing Workshop*. <https://doi.org/10.1109/MDSP.1989.97051>
- Dovis, F. (2015). *GNSS Interference Threats and Countermeasures*. Artech House GNSS technology and applications series. <https://us.artechhouse.com/GNSS-Interference-Threats-Countermeasures-P1710.aspx>
- Gao, G. (2007). DME/TACAN interference and its mitigation in L5/E5 bands. *Proc. of the 20th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2007)*. Fort Worth, TX, 1191–1200. <https://www.ion.org/publications/abstract.cfm?articleID=7623>
- García-Molina, J. A. & Fernández-Rubio, J. A. (2019). Array processing and unambiguous positioning of signals with multi-peak correlations. *Proc. of the 32nd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2019)*. Miami, FL, 3754–3758. <https://www.ion.org/publications/abstract.cfm?articleID=17024>
- Garzia, F., van der Merwe, J. R., Rügamer, A. & Felber, W. (2020a). Hardware implementation and evaluation of the HDDM. *2020 International Conference on Localization and GNSS (ICL-GNSS)*. Tampere, Finland. <https://doi.org/10.1109/ICL-GNSS49876.2020.9115507>
- Garzia, F., van der Merwe, J. R., Rügamer, A., Urquijo, S., & Felber, W. (2020b). HDDM hardware evaluation for robust interference mitigation. *Sensors*, 20(22), 6492. Retrieved from <https://doi.org/10.3390/s20226492>
- Garzia, F., van der Merwe, J. R., Urquijo, S., Taschke, S., Rügamer, A. & Felber, W. (2021). Sub-band agc-based interference mitigation. *2021 International Conference on Localization and GNSS (ICL-GNSS)*. Tampere, Finland. <https://doi.org/10.1109/ICL-GNSS51451.2021.9452250>
- Guerrero, J. M. & Gunawardena, S. (2017). Characterization of timing and pseudorange biases due to GNSS front-end filters by type, temperature, and Doppler frequency. *Proc. of the 2017 International Technical Meeting of The Institute of Navigation*. Monterey, CA, 418–444. <https://doi.org/10.33012/2017.14911>
- Ifeachor, E. C., & Jerjis, B. W. (2002). Digital signal processing: A practical approach, second edition. *Electronic systems engineering series*.
- Lathi, B. P., & Ding, Z. (2010). *Modern digital and analog communication systems, 4th edition*. Oxford series in electrical and computer engineering.
- Liu, A., An, J. & Wang, A. (2010). Effect of transform domain interference suppression on PN code tracking loops. *2010 IEEE 12th International Conference on Communication Technology (ICCT)*. Nanjing, China, 1064–1067. <https://doi.org/10.1109/ICCT.2010.5688540>
- Liu, Y., Yang, Y., Chen, L., Pan, H., & Ran, Y. (2020). Analysis of phase bias between GNSS signal components caused by nonideal group delay. *NAVIGATION*, 67(2), 291–305. Retrieved from <https://doi.org/10.1002/navi.363>
- Lucas-Sabola, V., Seco-Granados, G., López-Salcedo, J. A., Garcia-Molina, J. A. & Crisci, M. (2016). Cloud GNSS receivers: New advanced applications made possible. *2016 International Conference on Localization and GNSS (ICL-GNSS)*. Barcelona, Spain, 1–6. <https://doi.org/10.1109/ICL-GNSS.2016.7533852>
- Mallat, S. G. (1989). A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7), 674–693. <https://doi.org/10.1109/34.192463>
- Musumeci, L. & Dovis, F. (2012). A comparison of transformed-domain techniques for pulsed interference removal on GNSS signals. *2012 International Conference on Localization and GNSS*. Starnberg, Germany, 1–6. <https://doi.org/10.1109/ICL-GNSS.2012.6253131>
- O'Brien, A. J., & Gupta, I. J. (2010). An optimal adaptive filtering algorithm with zero antenna-induced bias for GNSS antenna arrays. *NAVIGATION*, 57(2), 87–100. <https://www.ion.org/publications/abstract.cfm?articleID=102520>
- Ojeda, O. A. Y., Grajal, J., & Lopez-Risueno, G. (2013). Analytical performance of GNSS receivers using interference mitigation techniques. *IEEE Transactions on Aerospace and Electronic Systems*, 49(2), 885–906. <https://doi.org/10.1109/TAES.2013.6494387>
- Pany, T. (2010). *Navigation signal processing for GNSS software receivers*. GNSS technology and applications series.
- Paonni, M., Curran, J. T., Bavaro, M. & Fortuny-Guasch, J. (2014). GNSS meta signals: Coherently composite processing of multiple GNSS signals. *Proc. of the 27th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2014)*. Tampa, FL, 2592–2601. <https://www.ion.org/publications/abstract.cfm?articleID=12322>
- Paonni, M., Jang, J. G., Eissfeller, B., Wallner, S., Rodriguez, J. A. A., Samson, J. & Fernandez, F. A. (2010). Innovative interference mitigation approaches: Analytical analysis, implementation, and validation. *2010 5th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing (NAVITEC)*. Noordwijk, Netherlands, 1–8. <https://doi.org/10.1109/NAVITEC.2010.5708055>
- Phelts, R., Walter, T. & Enge, P. (2009). Characterizing nominal analog signal deformation on GNSS signals. *Proc. of the 22nd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2009)*. Savannah, GA, 1343–1350. <https://www.ion.org/publications/abstract.cfm?articleID=8543>
- Popugaev, A. E., Wansch, R. & Urquijo, S. F. (2007). A novel high performance antenna for GNSS applications. *2nd European Conference on Antennas and Propagation (EuCAP 2007)*. Edinburgh, UK, 1–5. <https://doi.org/10.1049/ic.2007.1638>
- Querol, J., Onrubia, R., Alonso-Arroyo, A., Pascual, D., Park, H., & Camps, A. (2017). Performance assessment of time-frequency RFI mitigation techniques in microwave radiometry. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 10(7), 3096–3106. <https://doi.org/10.1109/JSTARS.2017.2654541>
- Raimondi, M., Julien, O., Macabiau, C. & Bastide, F. (2006). Mitigating pulsed interference using frequency domain adaptive filtering. *Proc. of the 19th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2006)*. Fort Worth, TX, 2251–2260. <https://www.ion.org/publications/abstract.cfm?articleID=6959>
- Rappaport, T. S. (Prentice Hall; 1996). *Wireless communications: principles and practice*.
- Rubino, D., Rügamer, A., Lukčín, I., Taschke, S., Stahl, M. & Felber, W. (2016). Galileo PRS snapshot receiver with server-side

- positioning and time verification. *Proc. of Symposium Positionierung und Navigation für intelligente Verkehrssysteme (POSNAV)*. <http://publica.fraunhofer.de/dokumente/N-464228.html>
- Rügamer, A., Rubino, D., Lukčič, I., Taschke, S., Stahl, M. & Felber, W. (2016). Secure position and time information by server side PRS snapshot processing. *Proc. of the 29th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2016)*. Portland, OR, 3002–3017. <https://doi.org/10.33012/2016.14781>
- Rügamer, A., van der Merwe, J. R., Cortes, I. & Felber, W. (2020). Theoretical and practical evaluation of an overlay multi-band front-end. *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. Portland, OR, 1160–1167. <https://www.ion.org/publications/abstract.cfm?articleID=17423>
- Senmoto, S., & Childers, D. G. (1972). Adaptive decomposition of a composite signal of identical unknown wavelets in noise. *IEEE Transactions on Systems, Man, and Cybernetics, SMC-2*(1), 59–66. <https://doi.org/10.1109/TSMC.1972.5408557>
- Smith, A. J. (1994). The need for measured data in computer system performance analysis or garbage in, garbage out. *Proc. Eighteenth Annual International Computer Software and Applications Conference (COMPSAC 94)*. Taipei, Taiwan, 426–431. <https://doi.org/10.1109/CMPSAC.1994.342768>
- Szumski, A. (2011). Finding the interference: Karhunen-Loeve transform as an instrument to detect weak RF signals. *Inside GNSS*, 56–64. <https://insidegnss.com/wp-content/uploads/2018/01/mayjune11-WP.pdf>
- van der Merwe, J. R., Garzia, F., Rügamer, A. & Felber, W. (2020a). High-rate DFT-based data manipulator (HDDM) algorithm for effective interference mitigation. *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. Portland, OR, 596–605. <https://doi.org/10.1109/PLANS46316.2020.9109950>
- van der Merwe, J. R., Garzia, F., Saad, M., Kreh, B., Rügamer, A., Monroy Gonzalez Plata, R. & Felber, W. (2020b). Receiver bandwidth compression for multi-GNSS signal processing. *Proc. of the 33rd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2020)*, 3671–3685. <https://doi.org/10.33012/2020.17681>
- van der Merwe, J. R., Rügamer, A., Garzia, F. & Felber, W. (2019). Wavelet based adaptive notch filtering to mitigate COTS PPDs. *Proc. of the 32nd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2019)*. Miami, FL, 3285–3299. <https://doi.org/10.33012/2019.17068>
- van der Merwe, J. R., Rügamer, A., Garzia, F., Felber, W. & Wendel, J. (2018). Evaluation of mitigation methods against COTS PPDs. *2018 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. Monterey, CA, 920–930. <https://doi.org/10.1109/PLANS.2018.8373470>
- van der Merwe, J. R., Garzia, F., Rügamer, A., Vidal, I. C. & Felber, W. (2020c). Adaptive notch filtering against complex interference scenarios. *2020 European Navigation Conference (ENC)*. Dresden, Germany, 1–10. <https://doi.org/10.23919/ENC48637.2020.9317518>
- Wang, W., Guo, M. & Chen, J. B. (2014). A new narrowband interference mitigation algorithm based on adaptive wavelet packet decomposition. *2014 Fourth International Conference on Instrumentation and Measurement, Computer, Communication and Control*. Harbin, China, 6–11. <https://doi.org/10.1109/IMCCC.2014.10>
- Wendel, J., Rügamer, A., van der Merwe, J. R., Urquijo, S., Pfaffelhuber, D., Waelkens, A. & Rüegg, D. (2019). Impact of receiver front-end characteristics on high order BOC tracking. *Proc. of the 32nd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2019)*. Miami, FL, 3738–3753. <https://doi.org/10.33012/2019.17023>
- Wolovich, W. A. (1974). *Linear multivariable systems*. Applied mathematical sciences: Springer.
- Zhang, Y., Wu, H. & Gao, Y. (2013). Transform domain interference suppression in GPS/BD-2 receiver based on fractional Fourier transforms. *Proc. of the 26th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2013)*. Nashville, TN, 3456–3463.
- Zhou, J., & Do, M. N. (2005). In Papadakis, M., Laine, A. F., & Unser, M. A. (Eds.) *Multidimensional oversampled filter banks*. Proceedings Volume 5914, Wavelets XI International Society for Optics and Photonics, SPIE. 671–682. <https://doi.org/10.1117/12.618209>

How to cite this article: van der Merwe JR, Garzia F, Rügamer A, Felber, W. Advanced and versatile signal conditioning for GNSS receivers using the high-rate DFT-based data manipulator (HDDM). *NAVIGATION*. 2021;68:779–797. <https://doi.org/10.1002/navi.441>