

FPGA Implementation of a TVWS Up- and Downconverter Using Non-Power-of-Two FFT Modulated Filter Banks

Vianney Anis, Jincheng Guo, Stephan Weiss, and Louise H. Crockett

Department of Electronic & Electrical Engineering, University of Strathclyde, Glasgow, Scotland, EU
 {vianney.anis,stephan.weiss}@strath.ac.uk

Abstract—This paper addresses an oversampled filter bank (OSFB) approach to up- and down-convert any or all of the 40 channels in the United Kingdom’s TV white space (TVWS). We particularly consider the use of non-power-of-two fast Fourier transforms (FFTs), which provides a greater choice of design parameters over existing OSFB implementations. Using a field-programmable gate array (FPGA) software defined radio (SDR) platform, we compare two different 40-point FFT-based implementations of the system — one fully parallelised, one serialised — with an existing design using a radix-two 64-point FFT in terms of implementation cost and power consumption.

I. INTRODUCTION

Future TV white space (TVWS) transceivers will be permitted to access one or more selected channels within a given frequency band depending on the radio’s geographical location [1]. For example, in the UK the TVWS spectrum covers the ultra high frequency (UHF) range from 470–790 MHz with 40 channels, each with an 8 MHz bandwidth. Therefore, because of the need for frequency-agile transceivers with strict spectral mask requirements [2], we have previously opted for a filter-bank based multicarrier approach, which relies on an oversampled discrete Fourier transform (DFT) filter bank. This approach is numerically efficient, as it permits the potential up- and downconversion of all 40 channels at the cost of a single channel transceiver with only a small overhead [3].

The typical implementation of the DFT uses the Cooley–Tukey algorithm [4]; while it provides a very simple and efficient implementation on hardware, it is well-known that this algorithm is limited to DFT sizes that can be expressed as a power-of-two. Multi-carriers system such as orthogonal frequency-division multiplexing (OFDM) or DFT modulated filter-bank multi-carrier (FBMC) are then usually implemented using a fast Fourier transform (FFT) of size $N = 2^k, k \in \mathbb{N}$. In the case of a TVWS FBMC system, an implementation using a power-of-two FFT may be inefficient, as it requires the computation of more DFT coefficients than needed to cover 40 channels of the TV spectrum and furthermore it increases the sampling rate in the filter-bank and at the input of the radio frequency (RF) modulation and filtering system.

In this paper we study the impact of the DFT implementation on the resources usage and power consumption on a field-programmable gate array (FPGA) software defined radio (SDR) platform. We propose two 40-point mixed-radix designs, one of which is a fully parallelised design, while the

other focuses on reducing the FPGA area usage by serialising the processing in the FFT.

Thus, in Sec. II, we provide a brief overview over our TVWS transceiver design, while Sec. III focuses on the design of a 40-point FFT. Sec. IV addresses FPGA implementation, with results analysed in Sec. V.

II. TRANSCEIVER SYSTEM

This section outlines the transceiver radio front-end [3], which aims to up- and downconvert any or all of the 40 channels in the UK’s TVWS spectrum. The overall design idea is described in Sec. II-A, with details on the FFT-modulated FBMC provided in Sec. II-B.

A. Overall System Outline

In the FBMC transceiver, as summarised in [3] and shown in Fig. 1, the conversion from baseband to digital RF is performed in two stages. Seen from the receiving antenna, a first stage — stage 1 — converts between an RF signals and a lower frequency intermediate signal whose rate enables it to be handled by a FPGA. Stage 2 is responsible for the multiplexing of the 40 TVWS channels into a single baseband signal in the transmitter (Tx), and the demultiplexing from the equivalent single baseband signal in the receiver (Rx) back into the 40 TVWS channels. This multiplexing and demultiplexing is performed by an oversampled filter bank-based multicarrier system.

In this work we will focus on the design of stage 2 as typical SDR platforms implement stage 1 in the RF daughter-board.




Fig. 1. Two-stage TVWS filter bank Tx (above) and Rx (below) as proposed in [3] with a FBMC system in stage 2.




Fig. 2. Stage 2 prototype filter with passband width of 8MHz.

B. Stage Two: Filter Bank-Based Multiplexer

The conversion between the 40 TVWS channels and the baseband signal required for stage 1 is performed with the help of an oversampled DFT-modulated filter bank with $K_2^{(i)}$ channels, operating as an FBMC transmultiplexer. The design is based on a 8 MHz-wide prototype as characterised in Fig. 2, whose transition band depends on the oversampling ratio. In our design, we have decided to sample the TVWS channels at 16 MHz, i.e. they are oversampled a factor of 2. This provides a sufficient transition band, to protect adjacent channels as well as keeping the prototype filter as short as possible.

The prototype filter is modulated by a DFT to the $K_2^{(i)}$ different band positions, which in the Rx operate as band selection filters to extract bandlimited TVWS channels which subsequently can be decimated by $K_2^{(i)}/2$. In the Tx, these filters follow an expansion by $K_2^{(i)}/2$ and fulfil the purpose of interpolation filters.

An efficient polyphase representation of the FBMC blocks ensures that the filtering is always operated at the lower rate [5], [6]. Further, a DFT filter bank enables a factorisation into a polyphase network consisting of operations that only involve the real-valued prototype filter, and a $K_2^{(i)}$ -point DFT [7]. As a result, the FBMC implementation for $K_2^{(i)}$ channels is just as costly as the conversion of a single channel, plus the cost of a FFT operation.

III. NON POWER-OF-TWO FFT

A. Motivation and Rationale

For the designs in [8], it was found that a higher decimation or expansion in stage 1 — and therefore a lower decimation or expansion in stage 2 — leads to a more numerically efficient design than vice versa. Nonetheless, the FPGA implementation in [3] ignored this optimality, as it had to rely on power-of-two FFTs inside the FBMC system as a design constraint. With $N = 2^k \geq 40$, $k \in \mathbb{N}$, the smallest possible number of channels covering the TVWS spectrum is $N = 64$, thus involving 24 unused channel that either need to be zero-padded in the transmitter or discarded in the receiver.

In order to attain a numerically more efficient design, for a hardware implementation we either need to rely on a 40 point DFT, or opt for a mixed-radix architecture, factorising $N = 40$ in $N = 40 = 2^3 \times 5$, i.e. building a 40-point FFT from a number of 8-point and 5-point FFTs. While the 8-point section can again rely on readily available power-of-two FFT blocks, the radix-5 block requires an explicit realisation [9], [10].




Fig. 3. Radix-5 FFT, as presented in [12]

B. Radix-5 FFT Block

A radix- N FFT can be derived from an N -point DFT. Given time domain coefficients x_n , $n = 0, \dots, (N-1)$, N Fourier coefficients X_k , $k = 0, \dots, (N-1)$ are calculated via evaluation of $X(e^{j\Omega}) = \sum_{n=0}^{N-1} x_n e^{-j\Omega n}$ at sample points $X_k = X(e^{j\Omega_k})$, $\Omega_k = 2\pi k/N$. With time domain samples x_n in a vector $\mathbf{x} \in \mathbb{C}^N$, and the Fourier coefficients in a vector $\mathbf{X} \in \mathbb{C}^N$, the DFT can be written as $\mathbf{X} = \mathbf{T}\mathbf{x}$. Defining the twiddle factor $W_N^k = e^{j2\pi k/N}$, the DFT matrix for an example of $N = 5$ is

$$\mathbf{T} = \begin{bmatrix} W_5^0 & W_5^0 & W_5^0 & W_5^0 & W_5^0 \\ W_5^0 & W_5^1 & W_5^2 & W_5^3 & W_5^4 \\ W_5^0 & W_5^2 & W_5^4 & W_5^6 & W_5^8 \\ W_5^0 & W_5^3 & W_5^6 & W_5^9 & W_5^{12} \\ W_5^0 & W_5^4 & W_5^8 & W_5^{12} & W_5^{16} \end{bmatrix}. \quad (1)$$

Exploiting the periodicity of the complex exponential, whereby $W_N^k = W_N^{k \pmod{N}}$, the matrix operation can be restructured using e.g. the Winograd approach [11] to yield a reduction in implementation cost over the 25 complex-valued multiplications and additions required for (1). Various solutions have been presented in the literature [9], [10], [12], which generally differ by small trade-offs between the number of adders and multipliers.

The radix-5 flow graph of [12] is depicted in Fig. 3, and selected here for implementation due to its low number of complex valued operations. In its presented form, this radix-5 FFT block requires 8 multipliers, to perform 4 real-valued gains on complex-valued signals and 36 adders for 18 additions of complex-valued signals.

C. 40 Point FFT

Using the radix-5 FFT of Sec. III-B, a mixed-radix implementation with a standard 8-point FFT can yield the desired 40 point FFT. This requires two stages, in which either five 8-point FFTs are followed by a reorganisation and eight radix-5 FFTs, or vice versa. For the latter organisation, the flow graph is shown in Fig. 4.




Fig. 4. 40 point mixed-radix FFT

D. Complexity

Without consideration for any overheads for re-organisation –parallel sample reorganisation comes at a very low cost in comparison to computation as long as the FPGA is not overcrowded– in Fig. 4, computational complexity of a radix-5 FFT was stated in Sec. III-B. With an 8 point FFT requiring 24 complex valued multiply-accumulate operations, i.e. 96 real valued multiplications, and 48 real valued additions, the overall complexity of a 40-point FFT in terms of real valued operations can be stated as

$$C_{\text{multiply}} = 8 \cdot 8 + 5 \cdot 96 = 544 \quad (2)$$

$$C_{\text{add}} = 8 \cdot 36 + 5 \cdot 48 = 528. \quad (3)$$

This can be contrasted to the 40^2 complex multiply accumulate operations, if the transform was implemented by a standard DFT, requiring 6400 real-valued multiplication and 3200 real-valued additions. Therefore, the mixed-radix FFT approach offers a reduction in computational complexity by one order of magnitude.

IV. IMPLEMENTATION AND DESIGN CONSIDERATIONS

This section gives an overview of the design environment as well as descriptions and justifications of various design decisions.

A. Platform and Experimental Setup

In this work we simulate a TVWS transceiver using a Xilinx FPGA-based SDR platform, composed of a ZC706 FPGA evaluation board and a RF daughter-board AD-FMCOMMS4. While this SDR system cannot cover the full TVWS spectrum, due to bandwidth limitation at the daughterboard, it is a typical setup for SDR, which shares its architecture with systems capable of higher performances such as the USRP N310; furthermore, our system should be representative of a TVWS SDR transceiver, and could be easily scaled up using a high performance FPGA and RF daughterboard.

To reduce the system complexity and guaranty convergence of the synthesis and implementation algorithms, we only implemented the FBMC system without all the adjacent subsystems required in a real-life scenario, such as synchronisation and equalisation processes. All systems were designed using a Simulink model and then implemented using the SDR Hardware/Software co-design workflow from Matlab.

B. Word-Length Considerations

In [3], it has been established that in order to keep the out-of-band emissions to adjacent TVWS channels below the -69dB currently suggested by the regulator [2], a word length of 12 bit must be used at RF for the filters designed for that purpose. If higher suppression is required, we can accomplish that with longer, more frequency-selective channels. Incorporating the resolution gain in the up- and downconversion stages, samples and coefficients at baseband should be resolved with 16 bits. But while these parameters are important, the wordlength limiting factors reside in the hardware, with the DSP48E1 block of a Xilinx FPGA discouraging wordlengths above 18 bits [13], in addition most digital-to-analogue converter (DAC) used in SDR platforms are limited to 14 bits, we then used 18 bits as a generic wordlength across the all system.

C. Serialisation of the Filters

Even in an efficient polyphase implementation, the prototype filters are of a length that requires a significant amount of multiplication-accumulation (MAC) operations, most likely exceeding the DSP48E1 resources available on the FPGA, if the filter bank is to be implemented using a fully parallel structure. While our work for the polyphase filter of the FBMC system is based on the design presented in [7], we decided to serialise as much as possible the multiplication operation to limit the number of DSP48E1s required. For simplicity of design –avoiding the introduction of various delays in the different branches of the polyphase architecture– we chose a serialisation factor K defined such that the FFT length N can be expressed as $N = mK, m \in \mathbb{N}$.

As we are operating at a scaled down frequency and a relatively small FPGA we use $K = N$, but in a real system a factor $K = N/m, m \in \{2, 4, 8\}$ would be required to operate at a reasonable frequency for a FPGA.

D. Transform Structure

We designed two versions of the non-power-of-two FFT, a fully parallel version as shown in Fig. 4 and a serial structure presented in Fig. 5.

The 40-point FFT serial version uses a 5-point FFT block as first stage operating at eight times the input sampling frequency F_s and 8-point FFT running at five times the input sampling frequency as second stage. This design then reduces the required multiplications by a factor of eight for the radix-5 stage and respectively five for the 8-point FFT stage. We should however see an increase in the use of other resources such as register, random access memory (RAM) and look-up table (LUT) which will be required for the serial-to-parallel




Fig. 5. Serialised 40 point mixed-radix FFT

and parallel-to-serial operation. Similarly the 48- and 56-point FFT use respectively a radix-3 and radix-7 as first stages combined with a 16- and 8-point second stage.

V. IMPLEMENTATION RESULTS

In this section, we present resource cost and power consumption of the systems described in Sec. III and Sec. IV.

A. Footprint

The resources used on the FPGA are presented in Tab. I; "FBMC N Parallel" –where $N = \{40, 48, 56\}$ – designates systems implemented using a N -point FFT with a fully parallel architecture and "FBMC N Serial" the serialised version of the transform as presented in Sec. IV-D, while "FBMC 64" is the FBMC system reconfigured to use a 64-point FFT block provided by Simulink.

As discussed in previous work [14], for signal processing applications the most critical resources on a FPGA are the MAC units, contained in the DSP48E1 modules of Xilinx's FPGAs. To that respect, the serialised versions of the FFT as proved more area-efficient for every implementation, allowing the larger transforms, such as the the 56- and 64-point FFT to fit on the FPGA, when a parallel version would use more resources than available. The reduced use of DSP48E1 however comes to the price of an increased number of RAM blocks necessary, which are most likely mapped to serial-to-parallel and parallel-to-serial conversions. This trade-off is however very advantageous as it make use of an otherwise unused resource, allowing to reassign the DSP48E1 block for

other tasks, such as equalisation and/or synchronisation [14]–[16].

The most area-efficient version of the system is the FBMC system implementing the 40-point serial version of the FFT, showing reduced block use against the 64-point FFT version in all type of resources.

B. Power Consumption

Power consumption data is produced using the implemented power report from Vivado and results are shown in Fig. 6, to help with clarity we only display the FPGA dynamic power consumption. The overall power consumption is obtained by adding the Zynq z7045 static overhead of 230mW, the ARM processor consumption of 1.57W and a 100mW off-chip consumption to the figures displayed in Fig. 6, adding 1.9W.

The results shown in Fig. 6 confirm the conclusions of Section V-A, as the 40-point FFT serial version proves to be the most energy efficient solution, with a 130mW saving compared to the previously designed 64-point version [3]. In our experiments, the area usage is the most critical variable when it comes to the power consumption. While the increase

resource available	LUTs	Flip-Flops	DSP48E1	RAM
FBMC 40 Parallel	117989 (53.97%)	218017 (49.87%)	632 (70.22%)	4 (0.73%)
FBMC 40 Serial	113621 (51.98%)	220783 (50.50%)	204 (22.67%)	18 (3.30%)
FBMC 48 Parallel	123033 (56.28%)	251615 (57.55%)	756 (84%)	4 (0.73%)
FBMC 48 Serial	128770 (58.91%)	244616 (55.95%)	336 (37.33%)	18 (3.30%)
FBMC 56 Serial	132220 (60.48%)	237517 (54.33%)	324 (36%)	18 (3.30%)
FBMC 64	125144 (57.25%)	229289 (52.44%)	629 (69.89%)	18 (3.30%)




Fig. 6. Power consumption results

of the working frequency leads to a higher power consumption from the clocks and the signals (transfers of data between blocks), the saving on the DSP48E1 power consumption largely counterbalances the increase.

One will note that the lower computational efficiency of the non power-of-two FFTs make 48-point FFT parallel and 56-point FFT serial version less energy-efficient, once again showing a strong correlation between the resources utilisation shown in Tab. I and the power consumption.

VI. CONCLUSIONS

In this paper we provided a new approach to the design and FPGA implementation of oversampled filter-bank multi-carrier systems for TVWS transmission, by moving away from the standard power-of-two FFT and considering a 40-point mixed-radix FFT. This approach has proven to be both less costly in terms of area and more energy-efficient by 230mW which represents a 6.7% energy saving for the overall transceiver compared to previous designs, when implemented on a Zynq z7045.

In a complete transceiver system, our approach might prove even more advantageous as systems upstream of the FBMC system in the receiver would run at a sampling frequency 30% lower than a 64-point FFT system.

ACKNOWLEDGEMENT

This work has received funding by the European Union Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 675891 (SCAVENGE).

REFERENCES

- [1] C. McGuire, M. R. Brew, F. Darbari, G. Bolton, A. McMahon, D. H. Crawford, S. Weiss, and R. W. Stewart, "HopScotch-a low-power renewable energy base station network for rural broadband access," *EURASIP Journal on Wireless Communications and Networking*, vol. 2012, no. 1, pp. 1–12, 2012.
- [2] S. J. Shellhammer, A. K. Sadek, and W. Zhang, "Technical challenges for cognitive radio in the TV white space spectrum," in *2009 Information Theory and Applications Workshop*, Feb. 2009, pp. 323–333.
- [3] R. A. Elliot, M. A. Enderwitz, K. Thompson, L. H. Crockett, S. Weiss, and R. W. Stewart, "Wideband TV White Space Transceiver Design and Implementation," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 63, no. 1, pp. 24–28, Jan. 2016.
- [4] J. W. Cooley and J. W. Tukey, "An Algorithm for the Machine Calculation of Complex Fourier Series," *Mathematics of Computation*, vol. 19, no. 90, pp. 297–301, 1965.
- [5] R. E. Crochiere and L. R. Rabiner, *Multirate Digital Signal Processing*. Englewood Cliffs, NJ: Prentice Hall, 1983.
- [6] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*. Englewood Cliffs: Prentice Hall, 1993.
- [7] S. Weiss and R. Stewart, "Fast implementation of oversampled modulated filter banks," *Electronics Letters*, vol. 36, no. 17, pp. 1502–1503, Aug. 2000.
- [8] R. A. Elliot, M. A. Enderwitz, F. Darbari, L. H. Crockett, S. Weiss, and R. W. Stewart, "Efficient TV white space filter bank transceiver," in *20th European Signal Processing Conference*, Aug. 2012, pp. 1079–1083.
- [9] F. Qureshi, M. Garrido, and O. Gustafsson, "Unified architecture for 2, 3, 4, 5, and 7-point DFTs based on Winograd Fourier transform algorithm," *Electronics Letters*, vol. 49, no. 5, pp. 348–349, Feb. 2013.
- [10] A. Karlsson, J. Sohl, and D. Liu, "Cost-efficient mapping of 3- and 5-point DFTs to general baseband processors," in *2015 IEEE International Conference on Digital Signal Processing*, Jul. 2015, pp. 780–784.
- [11] S. Winograd, "On computing the discrete Fourier transform," *Mathematics of Computation*, vol. 32, no. 141, pp. 175–199, 1978.
- [12] J. Löfgren and P. Nilsson, "On hardware implementation of radix 3 and radix 5 FFT kernels for LTE systems," in *2011 NORCHIP*, Nov. 2011, pp. 1–4.
- [13] R. Stewart and L. Crockett, "DSP for FPGA Primer," *University of Strathclyde. Glasgow*, 2011.
- [14] V. Anis, C. Delaosa, L. H. Crockett, and S. Weiss, "Energy-efficient implementation of a wideband transceiver system with per-band equalisation and synchronisation," in *2018 IEEE Wireless Communications and Networking Conference*, Barcelona, Spain, Apr. 2018, pp. 1–6.
- [15] S. Weiss, A. P. Millar, R. W. Stewart, and M. D. Macleod, "Performance of Transmultiplexers Based on Oversampled Filter Banks under Variable Oversampling Ratios," in *18th European Signal Processing Conference*, Aug. 2010, pp. 2181–2185.
- [16] S. Weiss, M. Harteneck, and R. Stewart, "On implementation and design of filter banks for subband adaptive systems," in *IEE Colloquium on Digital Filters: An Enabling Technology*, vol. 1998/252, Apr. 1998, pp. 12/1–12/8.