

WestminsterResearch

<http://www.westminster.ac.uk/westminsterresearch>

**Efficient Digital Signal Processing Techniques and Architectures
for On-Board Processors**

**Coskun, A., Kale, I., Morling, R.C.S., Hughes, R., Brown, S. and
Angeletti, P.**

This is an electronic version of a paper presented at the *3rd ESA Workshop on Advanced Flexible Telecom Payloads*, Noordwijk, The Netherlands, 21 to 24 March 2016.

Full details of the conference are available at:

<http://old.esaconferencebureau.com/2016-events/16c05/introduction>

The WestminsterResearch online digital archive at the University of Westminster aims to make the research output of the University available to a wider audience. Copyright and Moral Rights remain with the authors and/or copyright owners.

Whilst further distribution of specific materials from within this archive is forbidden, you may freely distribute the URL of WestminsterResearch: (<http://westminsterresearch.wmin.ac.uk/>).

In case of abuse or copyright appearing without permission e-mail repository@westminster.ac.uk

Efficient Digital Signal Processing Techniques and Architectures for On-Board Processors

3rd ESA Workshop on Advanced Flexible Telecom Payloads

21-24 March 2016

ESA/ESTEC
Noordwijk, The Netherlands

Adem Coskun⁽¹⁾, Izzet Kale⁽¹⁾, Richard C. S. Morling⁽¹⁾, Robert Hughes⁽²⁾, Stephen Brown⁽²⁾, Piero Angeletti⁽³⁾

⁽¹⁾*Applied DSP and VLSI Research Group, Department of Engineering,
University of Westminster, London, W1W 6UW, United Kingdom
Email: a.coskun@westminster.com, kalei@westminster.ac.uk, d.morling@westminster.ac.uk*

⁽²⁾*Airbus Defence and Space,
Gunnels Wood Road, Stevenage, Hertfordshire, SG1 2AS, United Kingdom
Email: robert.ro.hughes@airbus.com, stephen.brown@airbus.com*

⁽³⁾*European Space Agency,
Keplerlaan 1 Noordwijk, The Netherlands
Email: piero.angeletti@esa.int*

INTRODUCTION

Rapidly improving technologies make Digital Signal Processing (DSP) an increasingly cost efficient and highly reliable approach to enhance the flexibility and capacity of satellite communication links but carefully optimised designs are key to obtaining a power-efficient solution. Whether regenerative or transparent, processing communication payloads accommodate digital On-Board Processors (OBP) to perform advanced DSP techniques including digital beamforming, encoding/decoding, frequency conversion, routing, modulation/demodulation and digital channelisation. The OBP payload can be considered as the primary component of a flexible communication satellite. DSP can bring in more cost effective signal processing solutions than its analogue counterparts in terms of flexibility, scalability, architectural complexity and energy efficiency. Recently, digital OBPs are finding wide use in transparent satellite telecommunications to provide additional flexibility in frequency planning and routing for multi-beam systems. Modern systems for Mobile Satellite Services (MSS) using the lower frequency bands, where the available spectrum is heavily congested, are now almost exclusively implemented via digital OBPs but digital payloads are also increasingly becoming of interest for communications at higher frequency bands as the processing technology becomes more advanced and efficient, be it a narrowband or a broadband communication link, the OBP splits the received uplink signals into a large number of frequency channels which can be independently routed and further processed.

In this paper, we present a number of algorithmic and architectural DSP solutions to be incorporated in digital OBPs for communication satellites to boost the system performance primarily in terms of reducing their power consumption. More specifically this article addresses (1) Infinite impulse response (IIR) implementation of digital filters, (2) Efficiency savings in channeliser FFT twiddle storage and multiplications and their reconfigurable implementation (3) Companding of interconnect data, and (4) Critically sampled/reduced over-sampling channelisation. The applicability and efficiency of these approaches were evaluated in detail during our European Space Agency (ESA) funded research project entitled "Efficient Techniques for On-Board Processing", undertaken by Airbus Defence and Space and the Applied DSP and VLSI Research Group at the University of Westminster. The results demonstrated noteworthy improvements both in terms of power dissipation, and furthermore in the reduction of circuit complexity for future digital OBPs, which will be shown at the summary of results section.

EFFICIENT DSP TECHNIQUES AND ARCHITECTURES

Digital channelisation is a powerful approach to decompose the incoming signal into user channels. The Discrete Fourier Transform (DFT), as well as the Discrete Cosine Transform (DCT) filter banks, and Tree-structured filter banks are known to provide perfect reconstruction of the bandwidth of interest. For a DFT-based channeliser, both Digital Filter Bank (DFB) and the DFT Processing units cooperatively undertake the demultiplexing/multiplexing of user channels (In other fast transform based channelisation approaches, the same cooperation is present except the DFT being replaced by another fast transform).

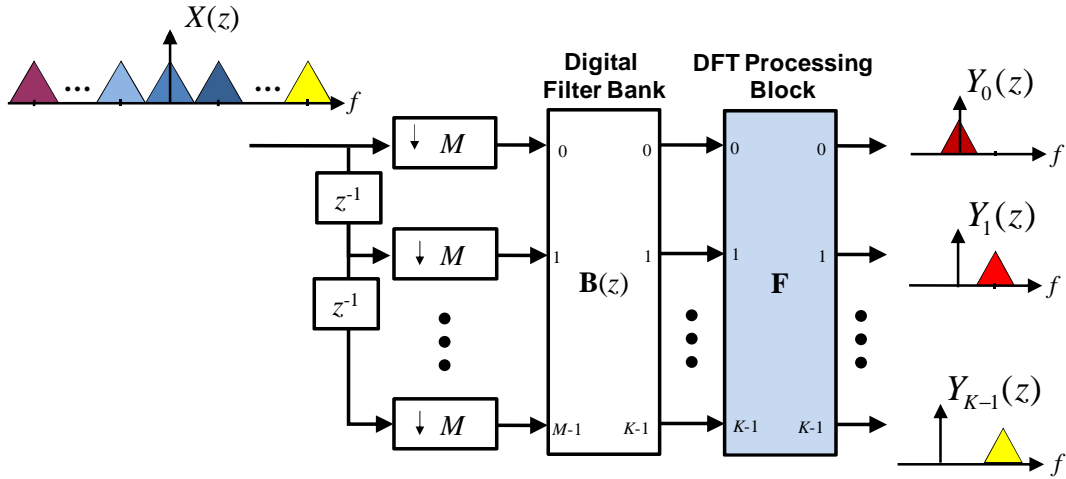


Fig. 1. DFT Polyphase DFB is Composed of (1) Unit Delayers and Down-samplers (2) DFB (3) DFT Processing Block.

To elaborate the necessity for these two components in a DFT-based channeliser (implemented in the form of a DFT polyphase DFB), Fig. 1 is given, where $X(z)$ is the input and $Y_i(z)$ for $i=1, \dots, K$ is the output for the DFT polyphase DFB. Please note that this system serves K users and that's why it can also be seen as a single input multiple output system as can be observed in Fig. 1. This structure is composed of the unit delayers, down-samplers, DFB and DFT Processing Blocks as labelled accordingly. The internal configuration of the DFB and the DFT Processing blocks can be formulated using two matrices; $\mathbf{B}(z)$ and \mathbf{F} respectively. While \mathbf{F} is the well known DFT matrix being composed of twiddle factors, the content and the dimensions of $\mathbf{B}(z)$ depend on the choice of K and M , which are the number of channels and the downsampling factor respectively [1].

In the following sub-sections we will detail some of the DSP approaches we identified in our earlier ESA funded project activity that DSP designers on the whole mostly ignore or overlook. These DSP approaches can effectively contribute to the efficiency of the OBP processor in terms of flexibility, power consumption and complexity.

Infinite Impulse Response (IIR) Implementation of Digital Filters

Irrespective of the implementation, DFT-based or otherwise, the essential functions are to divide the processed band into sub-channels and recombine them back for their retransmission. For both narrowband and broadband communication scenarios it is advantageous to implement these digital filters efficiently in the form of a polyphase Digital Filter Bank (DFB) as in Fig.1. The order, type and the structural arrangement of these DFBs depend on the channelisation requirements (such as the number of user channels, channel bandwidth and centre frequencies) and strongly affects the OBP implementation. The design of the digital filters (the requirements on its passband ripple, passband width, transition band width, stopband attenuation, guardband requirements), decision on the filter type, filter structure and the calculation of the filter coefficients are important design tasks for the realisation of the DFB. On the other hand, digital IIR filters are commonly ignored due to linear-phase and flat magnitude response requirements for distortion-free filtering, which can easily be obtained via the use of Finite Impulse Response (FIR) filters. Although an almost flat magnitude response is more realistically achievable using IIR filters, they cannot deliver perfect linear phase responses. There are several ways to linearise the phase response of an IIR filter. The strategy is normally to make use of allpass IIR filters and to design the linear phase response by manipulating the filter coefficients. For the state-of-the-art channelisers we have tested half-band IIR filters (where conventionally half-band band FIR filters are used in the reconfigurable digital OBPs) making use of Lowpass-to-Lowpass Frequency Transformation (LLFT) [2] and Balanced Model Truncation (BMT) [3] techniques. The choice of phase linearised half-band IIR filters was shown to yield savings in excess of 35% in comparison to their FIR counterparts both in terms of area and power dissipation in our earlier studies [4] and demonstrated that IIR filters bring in several advantages over FIR filters when used in the OBPs with minimal distortion on the phase achieving a peak-to-peak phase ripple of less than 2° across the passband region of interest and comfortably within acceptable limits.

On the other hand, the DFT has an important role in demultiplexing Frequency Division Multiple Access (FDMA) signals. In a multi-standard receiver a reconfigurable DFT unit can be a solution to adapt the satellite link to various communication standards. In a reconfigurable channeliser the size of the DFT unit is correlated with the number of channels to be served. Should the number of channels to be served change, then the size of the signal processing unit performing the DFT operation should also adapt its size to meet the varying number of channels. Therefore, the reconfigurability concept can be realized with a DFT unit where its size is also adaptable. The reconfigurability on DFT units can be realized in two different ways: (1) Accommodating various processing blocks with smaller sizes of DFT units, and therefore using a variety of DFT sizes generated by reconfiguring the connections




Fig. 2. Cascaded DFT units of small sizes

in between individual processing units (2) Deploy a Reconfigurable DFT processing unit that changes its size upon user request. In the following section we will address the reconfigurable DFT units, and their structural/ architectural implementation options.

Designing a Reconfigurable DFT Processing Block

DFT processing block is another important component of the DFT polyphase DFB in Fig. 1. It should be noted that the DFT size, i.e. K , is equal to the number of user channels in a DFB approach. If the number of channels to be served is subject to a change, then the size of the signal processing unit that performs the DFT operation should also adapt its size to meet the new number of channels. Therefore, the reconfigurability concept can be realized in a DFT unit of adaptable size. A straightforward approach to enable DFT operations of reconfigurable sizes is to accommodate smaller sizes of DFT units within the processor. DFTs of any size can be generated by reconfiguring the connections in between these processing units. This is indeed hardware inefficient and the reconfigurability of the connections between each small sized DFT unit could be power inefficient too. Rather than having many small sized DFT blocks, reconfigurable DFT units should be designed, which size up on request. Keeping in mind that K is a large number, which is the case in most of DFT operations, it is a much more efficient approach to divide a K point DFT into smaller sizes of DFT units. The Prime Factor Algorithm (PFA) enables the use of K_1 - and K_2 - point DFTs to perform a K point DFT, where $K = K_1 \times K_2$, if K_1 and K_2 are co-prime. The Winograd Fourier Transform Algorithm (WFTA) is a good example of a PFA. The minimum number of multiplications is achieved using Winograd's approach. The multiplications within a Winograd DFT unit are either real or imaginary, which avoids the use of complex multipliers. DFT blocks, either implemented using the radices of powers-of-two or using the WFTA, can be cascaded to form a larger size of DFT as shown in Fig.2. If the individual DFT processing blocks are designed to be reconfigurable, different sizes of DFTs can be obtained with the structure as in Fig. 2. In Fig. 2, N modules have been deployed to effect the configuration each of which is controlled by a reconfiguration control input.

The basic structure of a DFT module is depicted in Fig. 3 [5]. As can be seen from Fig. 3, each DFT module is composed of four sectors. If the "Reconfigurable Module" is configured to perform a Radix-2 or Radix-4 operation, then this would necessitate twiddle multiplications, which are performed in the processing unit named "Twiddle Multiplication".

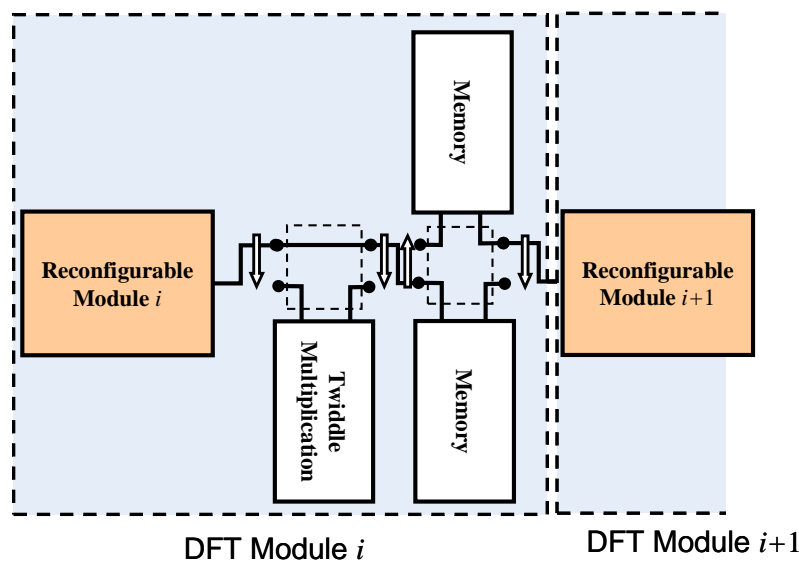


Fig. 3. The basic structure of a DFT module in Fig. 2

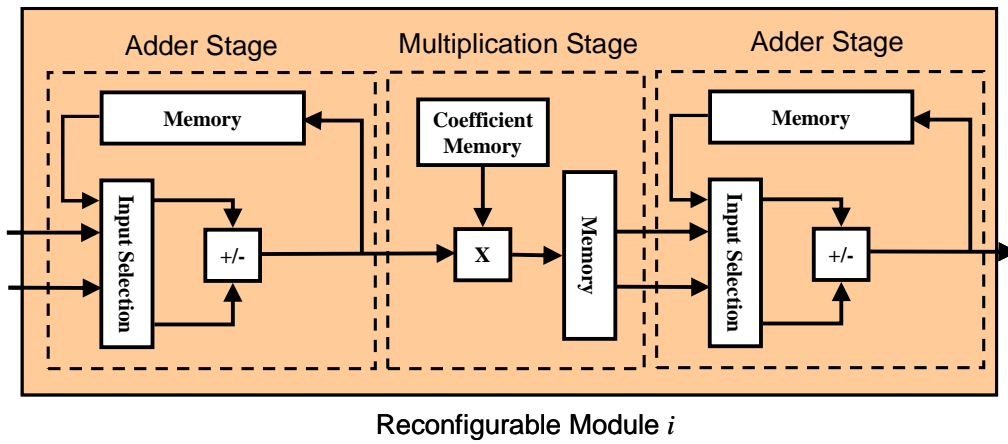


Fig. 4 Basic structure of a reconfigurable module in Fig. 3

On the other hand if the Reconfigurable Module is to perform a Winograd FFT, the twiddle multiplication stage should be avoided and the results from the reconfigurable module should be stored directly in the memories provided in the design. It should also be noted that for some of the DFT Modules the twiddle multiplication is necessary even if the WFTA is implemented. This is due to the different units not being mutually prime [5]. The reconfigurable module is shown in Fig.4 [6]. Due to the nature of the WFTA, the DFT should be performed in three stages; two adder stages and a real multiplication stage. Depending on the size of the DFT operation, the memory units set within the structure store the data and are accessed accordingly under the control of the input selection units. It is also possible to design a unified structure of small sized WFTA units to implement the reconfigurable module. Various DFT sizes such as 2-, 3-, 4-, 5-, and 7-point DFTs can be implemented in one unified structure [7] and a set of control signals can be used to select which DFT size to be utilised. There are possible extensions to this approach, one of which is pipelining. Our pipelined 5-point DFT approach presented in [8] (See Fig. 5.a), can be modified to be a reconfigurable DFT module. Using Reconfigurable Multiplier Blocks (ReMB) [9], the new structure has reduced complexity associated with the use of general purpose multipliers in the unified approach, making the structure more power and area efficient. The DFT Module given in Fig. 3 should be replaced by Fig. 5.b if the well-known Cooley and Tukey method is used utilising radices of powers-of-2. This necessitates a butterfly circuit for each DFT module and state store units to store the intermediate results while supplying input signals to the on-going butterfly operations as shown in Fig. 5.b. Twiddle multiplication is needed in some of the butterfly operations, where the twiddle factors can be generated using a twiddle factor generator or alternatively be fetched from a Read Only Memory (ROM).

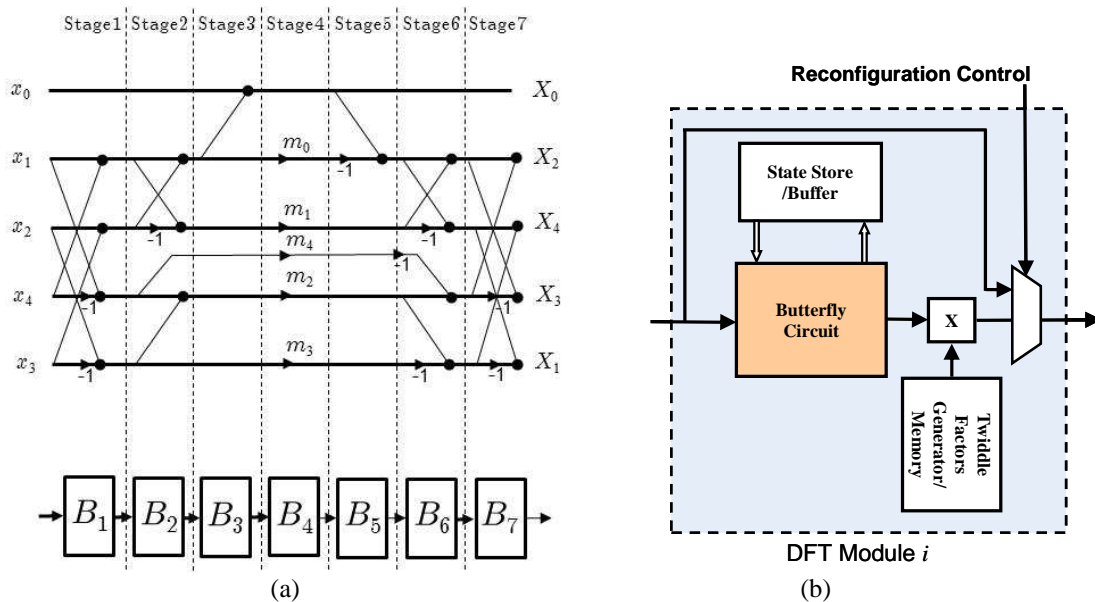


Fig. 5. (a) The plot on the top is the signal flow graph for a 5-point WFT and the plot on the bottom is the representation of our pipelined design with seven stages. (b) DFT modules for reconfigurability in mixed radix units

The two primary options to pipelined DFT butterfly structures are the Multipath Delay Commutator (MDC) and Single-path Delay Feedback (SDF). The SDF is a highly accepted structure over MDC as it can be implemented using a lower number of multipliers and memory units [10]. Radix-2, Radix-4, Radix- 2^2 , Radix- 2^3 and Radix- 2^4 are most widely used Radix units for the Radix Algorithm. The most important difference between Radix-2 and Radix-4 is that the butterfly circuits are fully utilised when performing Radix-4 while for Radix-2 they are utilised only for the half of the number of inputs. Radix- 2^2 has been shown to have lower number of multipliers and adders. It is also possible to have Radix- 2^3 and Radix- 2^4 but due to an increased complexity radices beyond 2^4 are regarded as being impractical [9]. It is possible to have a reconfigurable DFT module capable of implementing all Radices from Radix-2 up to Radix- 2^4 [11]. This sort of reconfigurability can be maintained by switching off (clock off or clock gating) the earliest DFT modules and feeding the data directly into the later stages. This is illustrated by the use of a multiplexer in Fig.5.b [11]. This type of reconfigurability makes it possible to have different sizes of DFT lengths of order 2 (i.e. 2-4-8-16-32-64-128-256...). As can be understood, a finer resolution is not possible using this type of reconfigurability with the Radix Algorithm. If a DFT size of 2048 is being used and due to the necessities on the communication link the DFT length is to be decreased, a length of 1024 is achieved by switching off one of the butterfly circuits. This type of reconfigurability is generally adopted in OFDM applications (e.g. Digital Subscriber Line (DSL), Digital Video Broadcasting (DVB) standards) where the OFDM sizes are in the orders of 2. However, for digital channelisation a finer resolution cannot be achieved using this approach, and therefore the structure in Fig.3 with a mixed Radix and Winograd modules seems to be the primary choice for the design of a reconfigurable DFT.

On the other hand, for a K -point DFT operation only $K/2$ twiddle factors are needed. However, it is not necessary to store all the twiddle factors in a DFT operation, and it is possible to store only $K/8$ twiddles and generate the rest from the stored values, which means only $K/4$ real numbers are needed. In our previous study, a number of Look Up Tables (LUTs) were used to store the twiddle factors and were replaced by a single LUT and an address generator was proposed to enable the use of this LUT for supporting the twiddle operations [12]. The storage strategies for the twiddle factors is another concept to bring in savings both in terms of circuit complexity and power consumption.

Critically Sampled / Reduced over-Sampling Channelisation

The applicability of polyphase quadrature mirror filters to a space frequency demultiplexing/re-multiplexing switching solution [13] was also investigated for On-Board processors. The specific interest in this method is that it appears to allow the transmission of channelised data critically-sampled between the demux and remux sections of the processor, which reduces the data rate requirement for the processor interconnect. This interconnect data rate is a major driver on the power dissipation and bandwidth capabilities for the switch network in high throughput transparent processors. As the processing technology performance improves, this “interconnect bottleneck” becomes progressively more important than the amount of arithmetic executed in the channelisation processing.

The principle of the algorithm is to divide the signal into multiple bands using a filter and decimate without loss of information. The algorithm then re-combines the signal using a complementary filter with arbitrarily small signal distortion. The filter properties required are:

- The aliasing components introduced by the decimating filter precisely cancel those generated by the multiplexing filter.
- The filter frequency responses seen at the channelised level overlap and add in such a way that the composite frequency response is nearly flat for all frequencies.
- These properties are simply achieved by constraining the filter response in the transition bands.

Such algorithms are typically used in sub-band coding applications when a wider bandwidth signal is decomposed into two or more sub-band signals which are subsequently recombined to reconstruct the original signal. This approach works well to support near-perfect reconstruction of wide channels from a multiplex of contiguous narrow sub-channels in a satellite communications application but fails at the boundaries between channels routed differently. The reason is that, where perfect reconstruction is not required, the expected cancellation between aliased sub-channels does not occur and generates spurious signals in the adjacent sub-channel. As a consequence it is necessary to route a “guard” sub-channel either side of each wide channel, lowering the spectral efficiency. To recover this efficiency it is necessary to construct a filter bank with twice as many channels as for the oversampled case, negating the advantage of the critical sampling in those applications requiring a large number of non-contiguously routed channels. For broadband scenarios in which a large number of sub-channels is used primarily to provide flexibility in bandwidth for a smaller number of wide channels, however, the savings in interconnect by using critical sampling can be significant. A modified algorithm was proposed whereby a conventional oversampled DFT-based filter bank can be converted to a critically sampled version using a tiny conversion stage. This allows a common design to be reconfigured to select the more appropriate choice between oversampled or critically sampled operation for different applications.

Table 1. Narrowband channeliser scenario summary

Function	Ref. Contribution	Reduction	Basis for reduction
Interconnect	42%	15%	2-bit companding
Digital Filtering/DFB	36%	35%	IIR design
DFT Processing	9%	16%	DFT using Reconfigurable Multiplier Blocks
Twiddles	11%	15%	Efficient Coefficient Store
Other	3%	0%	

Companding of Interconnect Data

The name "Companding" comes from the combination of "Compressing" + "Expanding" and refers to methods of reducing the data rate at particular points within a data processing algorithm. Since interconnection bandwidth is a major factor in the power consumption of the reference on-board processing architectures, companding can reduce that bandwidth while at the same time having a negligible impact on the end-to-end performance. The companding should be applied on the links between the ASICs as a transparent process that is not visible to the other processing algorithms within the ASICs.

Companding schemes based on entropy coding of the real and imaginary parts of each complex sample of sub-channelised data transmitted between ASICs has been investigated and a large number of coding schemes analysed to select the best performing ones. These schemes increase the precision used to represent small signals at the expense of lower precision and increased quantisation noise for large signals. Because the signal statistics are heavily weighted towards low amplitude values, this results in a net performance improvement, or data reduction for the same performance. Using coding schemes with very low implementation costs it is typically possible to reduce the interconnect wordlength by 2 bits with little loss of performance, resulting in a 15% saving in interconnect on 12-bit data words.

SUMMARY OF RESULTS

To evaluate the performance of each technique cited in this article a flight-representative reference OBP design has been chosen and individual savings contributed by each technique to the total power dissipation were calculated. We observed that between them they tackle all the main power contributions in the reference DSP architectures and most can be used in combination to deliver cumulative power savings. The quantitative overall savings depend heavily on the target technology and the system scenario. Technology parameters affecting this include the relative contributions of interconnect, RAM and combinatorial logic to the power dissipation. System parameters include the port bandwidth, the sub-channel granularity, the guard band width, the traffic capacity and the amount of digital beam-forming required. For the selected narrowband communication scenario without digital beam-forming, the main power contributions in the reference design are split between functions as shown in Table 1. The biggest savings are due to interconnect companding and the use of IIR filters. On the other hand the Reconfigurable DFT processing is currently under investigation by the authors and will be reported in a future article.

CONCLUSION

In this paper a number of DSP techniques and architectures are exposed that improve the efficiency and reconfigurability of the digital On-Board Processors (OBPs) for satellite communications. The main focus of this work was set to be the channelisation of the incoming spectrum using DFT polyphase Digital Filter Banks. The design of the filter bank making use of Infinite Impulse Response filters brought in noteworthy savings both in terms structural complexity as well as reducing the power dissipation of the channelisation unit within the digital OBP. The motivation here was to design IIR filters that can replace their FIR filter counterparts. It has been observed that IIR filter are capable of undertaking the DFB tasks causing only a minimal distortion in the phase response. On the other hand reconfigurability on the number of user channels and bandwidth can be achieved utilizing a reconfigurable DFT unit. In this paper we have reported on efficient DSP structures to deliver a reconfigurable DFT operation. Using Reconfigurable Multiplier Blocks (RMBs) the complexity of the reconfigurable DFT processor was reduced enabling a cost effective realisation. Furthermore, we have also investigated the companding and critical sampling techniques as well as efficient coefficient store strategies applicable to digital OBPs. How much each of these techniques contribute to the reduction in the overall power consumption within the processor has been shown in this article and it has been observed that the biggest savings are due to interconnect, companding and the use of IIR filters, while the other proposed techniques achieved moderate savings too.

REFERENCES

- [1] K. Eneman and M. Moonen, "DFT modulated filter bank design for oversampled subband systems," *Signal Process.*, vol. 81, pp. 1947–1973, 2001.
- [2] A. Krukowski and I. Kale, *DSP system design: Complexity reduced IIR filter implementation for practical applications*, Boston: Kluwer Academic Publishers, 2003
- [3] I. Kale, G. D. Cain, and R. C. S. Morling "Minimum-phase filter design from linear-phase startpoint via balanced model truncation", *IET Electronic Letters*, vol. 31, no 20, pp. 1728–1729, 1995.
- [4] A. Coskun, I. Kale, R.C.S. Morling, R. Hughes, S. Brown, and P. Angeletti, "Halfband IIR Filter Alternatives for On-Board Digital Channelisation", in *Proc. AIAA Int. Communications Satellite Syst. Conf.*, Florence, Italy, Oct. 2013.
- [5] F. Camarda, J.-C. Prevotet, and F. Nouvel, "Implementation of a reconfigurable fast Fourier transform application to digital terrestrial television broadcasting," in *Proc. Int. Workshop Field-Programmable Logic Applications*, pp. 353–358, 2009.
- [6] F. Camarda, J.-C. Prevotet, and F. Nouvel, "Towards a Reconfigurable FFT: Application to Digital Communication Systems", *Fourier Transforms, Theory and Applications*, p.185-202, InTech, 2011.
- [7] F. Qureshi, M. Garrido, O. Gustafsson, "Unified architecture for 2, 3, 4, 5, and 7-point DFTs based on Winograd Fourier transform algorithm", *IET Elect. Letters*, vol. 49, no. 5, pp. 348-349, Feb. 2013.
- [8] A. Coskun, I. Kale, R.C.S. Morling, R. Hughes, S. Brown, and P. Angeletti, "The Design of Low Complexity Low Power Pipelined Short Length Winograd Fourier Transforms," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, 2014, pp. 2001–2004.
- [9] S. S. Demirsoy, I. Kale, and A. G. Dempster, "Reconfigurable multiplier blocks: Structures, algorithm and applications," *Circuits, Syst. Signal Process.*, vol. 26, no. 6, pp. 793–827, Dec. 2007
- [10] C. Yang, T. Yu, and D. Marković, "Power and Area Minimization of Reconfigurable FFT Processors: A 3GPP-LTE Example", *IEEE Trans. Solid State Circuits*, vol.47, no.3, pp. 757-768, Mar. 2012.
- [11] Q. Lu, X. Wang, and J. Niu, "A low-power variable length FFT processor base on Radix-24 algorithm," In *Proc. Asia-Pacific Conference on Postgraduate Research in Microelectronics & Electronics*, pp.129-132, Jan. 2009.
- [12] A. Coskun, S. Cetinsel, I. Kale, R.C.S. Morling, R. Hughes, S. Brown, and P. Angeletti, "Efficient Coefficient Store in Decomposed DFT/FFT Architectures for On-Board Processors", in *Proc. AIAA Int. Communications Satellite Syst. Conf.*, Florence, Italy, Oct. 2013.
- [13] J. H. Rothweiler, "Polyphase Quadrature Filters – a new Sub-band coding technique", *IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP (1983)* pp1280-1283.