DTU Library

Multiplier-free filters for wideband SAR

Dall, Jørgen; Christensen, Erik Lintz

[Link back to DTU Orbit](#)

Multiplier-free Filters for Wideband SAR

Jørgen Dall and Erik Lintz Christensen

Ørsted•DTU, Technical University of Denmark
Ørsteds Plads, 348, DK-2800 Kgs. Lyngby, Denmark
Phone: +45 4525 3800 / Fax: +45 4593 1634 / e-mail: jd@oersted.dtu.dk

Abstract–This paper derives a set of parameters to be optimized when designing filters for digital demodulation and range prefiltering in SAR systems. Aiming at an implementation in field programmable gate arrays (FPGAs), an approach for the design of multiplier-free filters is outlined. Design results are presented in terms of filter complexity and performance. One filter has been coded in VHDL and preliminary results indicate that the filter can meet a 2 GHz input sample rate.

Index Terms–multiplier-free filters, filter optimization, digital demodulation, SAR.

I. INTRODUCTION

The filters considered in this paper are intended for the digital front end of a wideband SAR system like the SAR++ system, the next generation airborne SAR, under development in Denmark. Two types of filters are addressed: digital demodulation (DDEM) filters and range prefilters (RPF).

Technology now allows the analog-to-digital conversion in wideband SAR systems to be moved from baseband to intermediate frequency. Emerging analog-to-digital converters (ADCs) seem to offer sufficient fidelity for digital demodulation up to 1-2 GHz bandwidths [1], which makes them adequate for wideband SAR systems.

Implementing the quadrature demodulation in the digital domain has several advantages. The ADC transfer characteristic does not contribute to the channel mismatch, when only one ADC is involved. Perfect orthogonality is achieved since the mixers are digital multipliers and the local oscillator mixer input signals are digitally generated. The digital demodulator does not have low frequency or DC offset problems as the DC ends up outside the bandwidth of the signal.

Range prefilters are used to improve the signal-to-noise ratio when resolution is traded off for swath width. By sampling the backscattered pulses at the highest possible sampling rate the thermal noise and quantization noise is spread over the largest possible bandwidth. Thus, a lowpass filter with a pass band equal to the smaller signal bandwidth improves the signal-to-noise ratio. By combining the filtering with decimation, the computational requirements are reduced and the swath width increased for a given buffer size.

In wideband SAR systems both filter types make stringent hardware requirements. To be implemented in Field Programmable gate array (FPGA) technology the filters are desired to have few coefficients, each with only a few one-bits so that shift-and-add operations can substitute the multiplications by the signal samples.

Section II addresses the digital demodulator theory. Section III derives a set of filter parameters to be optimized in the filter design. A novel global search technique is used for the optimization, while Section IV is confined to a description of the search space. Finally, Section V gives examples of the complexity and performance of a few filters.

II. DEMODULATION FILTERS

The DDEM is implemented with a local oscillator frequency of $f_o = f_s/4$ and a phase chosen such that the I- and Q-channel mixing becomes multiplications by $\cos(2\pi n/4)$ and $-j\cdot\sin(2\pi n/4)$, respectively. f_s is the sampling frequency and n is the sample index. Thus, the mixing requires nothing but a multiplication of the I- and Q-channel samples by 1, 0, -1, 0… and 0, -1, 0, 1…, respectively, i.e. the samples are simply passed, reversed, or zeroed.

Fig. 1 illustrates that the mixing converts the real input signal with spectrum S into two signals with spectra

$$I'(f) = [S_+(f) + S_+(f - 2f_o)]/2 \tag{1}$$

$$Q'(f) = j[S_-(f) - S_-(f - 2f_o)]/2. \tag{2}$$

In terms of the spectra of the filtered and decimated signals

$$I(f) = S_+(f)/2 \tag{3}$$

$$Q(f) = j\,S_-(f)/2 \tag{4}$$

the desired output spectrum of the demodulator is

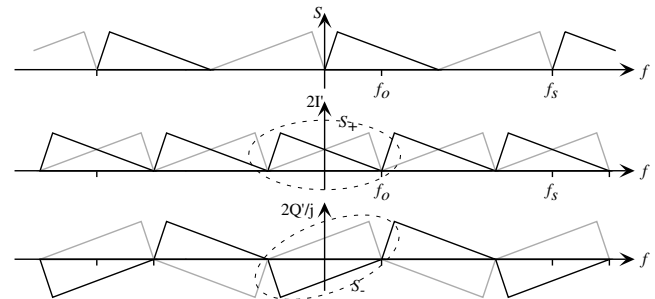$$S(f) = I(f) + j\,Q(f). \tag{5}$$

Fig. 1. Digital demodulator spectra. The gray spectra are the complex conjugate of the black ones.

S_{iq}, S_+, and S_- are spectra with maximum bandwidth $\pm f_o$. Hence, after the quadrature mixing an ideal lowpass filtering should in principle remove the I' and Q' components centered at $f = 2f_o$ and decimate by a factor of two. However, (1) shows that the same I-signal (multiplied by 2) is obtained by a direct decimation as the aliased spectral bands are identical. According to (2) the same technique can be applied to the Q-signal if the phase of the spectral component at $f = 2f_o$ is first changed by π. Otherwise the result become zero as the aliasing cancels the spectrum completely (the time domain explanation is that all the non-zero samples are rejected by the decimation).

The spectral phase shift of π means that the even samples are found and the odd samples are zeroed, thereby lining up the Q-samples with the I-samples. In the time domain this is accomplished by filtering the Q-signal with an ideal lowpass filter (passband $\pm f_o$) with zeroed even samples (all signals with zero-valued even samples have the form (2) just like all signal with zero-valued odd samples have the form (1)).

In practice the lowpass filtering (or interpolation) is not ideal and so the transfer functions of the two channels become different as in [1]. If the two channels are not combined in the SAR front end (e.g. no motion compensation or pulse compression) it is possible to equalize the Q-channel offline, where longer filters can be applied without jeopardizing real-time constraints. Otherwise, a channel balance improvement can be accomplished by shaping the I-spectrum online thus matching it better to the Q-spectrum. This approach allows an equalization to be applied off-line even if the two channels have already been combined.

For the interpolation (spectral phase shift) and spectral shaping, two decimation FIR filters are used

$$h_i(n) = \begin{cases} h(n), & n \text{ even} \\ 0, & n \text{ odd} \end{cases} \qquad (6)$$

$$h_q(n) = \begin{cases} 0, & n \text{ even} \\ h(n), & n \text{ odd}. \end{cases} \qquad (7)$$

$h = h_i + h_q$ is an approximation to an ideal lowpass filter with passband $\pm f_o$.

An alternative to h_i and h_q is interpolated shifts of the I- and Q-signals by $\pm(8f_o)^{-1}$, respectively. This approach has the advantage that the I- and Q-filters have the same amplitude spectrum, as the interpolators are mutually time reversed. However, the ideal interpolators have a phase jump of $\pi/2$ at f_o and hence the phases become non-linear in practice (the time domain explanation is that the interpolators are not symmetric). Offline phase equalization may be prevented by the fact that the I and Q phase nonlinearities are different, but most importantly, the filter implementation cannot profit from any filter symmetry.

III. PERFORMANCE PARAMETERS

The DDEM and RPF filters are optimized with respect to a set of quality parameters.

A. Demodulation filter

When equalization is applied, the often used sidelobe properties are not very sensitive quality measures. Instead h_i and h_q are Fourier transformed and the relative power spectrum difference is adopted as a measure of the channel imbalance

$$\Delta_{iq} = \frac{1}{B_s} \int_{B_s} \frac{\left|H_i(f)\right|^2 - \left|H_q(f)\right|^2}{\left|H_i(f)\right|^2} \, df. \qquad (8)$$

The integration is confined to the signal bandwidth B_s, disregarding a potential transition band, B_t. In the SAR++ case $B_s = 0.8 \cdot (2f_o)$ and $B_t = 0.2 \cdot (2f_o)$.

The signal-to-thermal-noise ratio is not changed by the filtering, equalization, and subsequent spectral weighting by $W(f)$ (input signal and noise spectra assumed flat). The signal-to-processing-noise ratio, however, might be jeopardized by a strong equalization. For each filter and for a given $W(f)$ the resulting processing noise can be reduced by minimizing

$$N_p = \frac{1}{B_s} \int_{B_s} \left|\frac{W(f)}{H(f)}\right|^2 df \cdot \frac{N_H}{N_W}, \qquad H \in \{H_i, H_q\} \qquad (9)$$

$$N_H = \frac{1}{B_s} \int_{B_s} \left|H(f)\right|^2 df \qquad (10)$$

$$N_W = \frac{1}{B_s} \int_{B_s} \left|W(f)\right|^2 df \qquad (11)$$

The processing noise (e.g. caused by rounding in the adder tree of the filter implementation) is assumed to be white when referred to the filter output. $1/|H(f)|^2$ accounts for the equalization and $|W(f)|^2$ for the weighting. N_H and N_W are normalization factors. Note that N_p is unity if H is constant or proportional to W, but larger if the equalization amplifies the edges of the spectrum much more than the weighting attenuates the same part of the spectrum.

B. Range prefilter

The primary figure of merit is the integrated noise energy that is aliased into the pass band from the $M-1$ stop bands

$$N_a = \frac{1}{B_s} \int_{B_s} \frac{\sum_{m=1}^{M-1} \left|H_m(f - mf_s/M)\right|^2}{\left|H(f)\right|^2} \frac{|W(f)|^2}{N_W} \, df. \qquad (12)$$

M is the decimation ratio. The average is computed over the signal bandwidth B_s while the sum is over the spectral bands aliased into B_s. The W/H ratio accounts for the equalization and weighting.

N_p, is also relevant to range prefilters, but less important than N_a. Since prefiltering calls for identical I- and Q-filters, Δ_{iq} is always unity (unless demodulation filtering and prefiltering are combined in one filter).
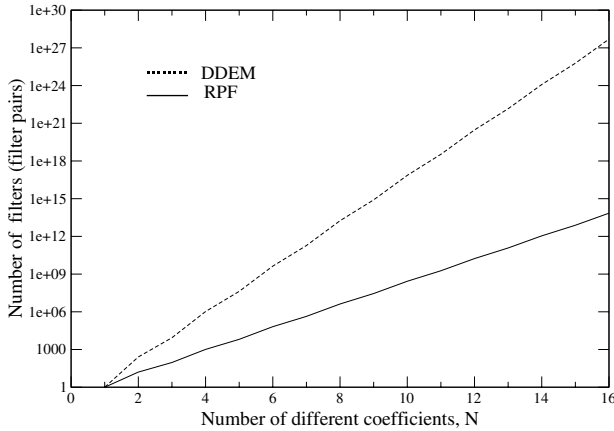
Fig. 2. Search space size. At most half of the 3 bit positions allocated per coefficient are used (average).

Fig. 3. DDEM channel imbalance $(|H_i(f)|^2 - |H_q(f)|^2)/|H_q(0)|^2$.

IV. FILTER OPTIMIZATION

A novel search algorithm is used for the filter optimization, but this algorithm is outside the scope of this paper.

The search space is minimized in various ways. The number of different coefficients to be determined, N, is (almost) halved by the fact that the filters are symmetric (all but the DDEM Q-filters have an odd number of coefficients). On the other hand the demodulation I- and Q-filters must be found in a joint optimization.

The coefficients are not complex. The word length is 8-bits, but when allocating positions for potential one-bits the filter coefficients are assumed to decrease at least as fast as those of an ideal lowpass filter. Also, a fixed power-of-two exponent is defined for each coefficient, typically leaving only three positions open for one-bits. The signs of certain coefficients are often known a priory. However, the signs of coefficients near expected filter nulls or near the ends of the filter must always be included in the optimization. Fig. 2 illustrates that even with these constraints the size of the search space is enormous.

V. RESULTS

Filters have been designed for demodulation and for prefiltering with decimation factors 2 and 4. Examples are given in Table 1 where the filter complexity is expressed in terms of the number of different coefficients, N, and the average number of bits per coefficient, b. The fact that N_p is smaller than 0 dB shows that the combined effect of equalization and

weighting does not amplify the processing noise. The N_a performance is adequate as no signal is present in the stop bands. The channel balance is so good that a common I- and Q-equalization yields a 3 dB pulse width that is just 0.001 % inferior to that achieved with ideal filters. The peak and integrated sidelobe ratios are inferior by -0.002 dB and 0.04 dB, respectively.

The uppermost demodulator in Table 1 has no I-filter (it is a delta-function) and its Q-filter is identical to filter number 3 in [1]. Fig. 3 compares the channel balance of the two demodulator filter pairs in the table.

Filters are being coded in VHDL for subsequent implementation in FPGA technology. Preliminary results indicate that the filters can operate at data rates up to 2 GHz.

VI. CONCLUSIONS

Ideally, no filtering should be applied to the I-channel of a digital quadrature demodulator. However, if the imperfect interpolation in the Q-channel calls for offline equalization, an I-channel filtering might be required too. Otherwise, pulse compression and other operations that combine the two channels cannot be done online.

Parameters for optimization of the channel balance and the signal-to-noise ratio have been presented and used to design multiplier-free filters for demodulation and range prefiltering. The low filter complexity allows FPGA implementations suitable for wideband SAR systems.

REFERENCES

[1] J.H. Jørgensen, "Digital Demodulator for Wide Bandwidth SAR", Proceedings of the International Geoscience and Remote Sensing Symposium, IGARSS'2000, Vol. V, pp. 2269-2271, July, 2000.

TABLE 1.
FILTER PERFORMANCE AND COMPLEXITY.

	N	b	Δ_{iq} [dB]	N_p [dB]	N_a [dB]
DDEM, I	1	1.0	-11.6	0.00	NA
DDEM, Q	4	1.5		-0.06	NA
DDEM, I	4	1.5	-17.4	-0.22	NA
DDEM, Q	4	1.5		-0.20	NA
RPF, M=2	10	1.5	NA	-0.21	-38.7
RPF, M=4	14	1.5	NA	-0.22	-31.3