DEFINING THE WORDLENGTH OF THE FRACTIONAL INTERVAL IN INTERPOLATION FILTERS

Francisco López,* Jussi Vesma** and Markku Renfors*
** Institute of Communications Engineering, Tampere University of Technology
P.O. Box 553, FIN-33101, Tampere, FINLAND
** Nokia Research Center, Helsinki, Finland
E-mail: {lopezm,mr}@cs.tut.fi; Jussi.Vesma@nokia.com

ABSTRACT

The fractional interval is used to determine the interval between the output sample and the previous input sample in interpolation filters. In this paper, the effects of the quantization of this parameter are studied and the wordlength is derived for different kind of interpolation filters having anti-aliasing or anti-imaging properties.

1. INTRODUCTION

Fractional interpolation filters or fractional delay filters are used in various digital signal processing applications to evaluate sample values of a discrete-time signal at arbitrary points between the existing samples [Laa96]. This paper concentrates on the interpolation filters having one continuous-valued input parameter μ_l which is called as the fractional interval. This parameter is used to determine the time interval between the interpolated output sample $y(l)$ and the previous input sample $x(n_l)$. In most of the applications it is required that the fractional interval is adjustable during the computation.

Historically, VLSI implementation of digital signal processing algorithms has required fixed-point arithmetic for the sake of cost, speed and power consumption, and it has been important to use the fewest number of bits possible to carry each signal in the system. The strong demand for wired/wireless communication services, among others, calls for low-cost high-volume IC solutions (ASIC or FPGA) that can perform real-time operations using as few chips as possible. Hence, specific designs that are to be implemented using ASIC or FPGA need to be optimised in terms of wordlength and architecture, in order to keep cost, size and power consumption to a minimum.

In [Lop00], the effects of quantizing the fractional interval have been studied in theory. In this paper, the effects of finite wordlength representation of μ are further examined, with special attention on the different requirements for different kind of interpolation filters in terms of number of bits.

This work was carried out in the project "Advanced Transciever Architetures and Implementations for Wireless Communications" supported by the Academy of Finland.

2. QUANTIZING THE FRACTIONAL INTERVAL

In order to analyse the finite precision effects of the fractional interval we shall make use of the hybrid analog/digital model, illustrated in Fig. 1. Any discrete-time interpolation filter can be obtained from a continuous-time filter $h_a(t)$, that is, there is an underlying continuous-time impulse response filter inherent to the interpolation process, satisfying the following condition [Ves99]

$$h(k,\mu(l)) = h_a\big((\mu(l)+k)T_s\big), \tag{1}$$

for $k = -N/2, -N/2+1 \ldots N/2-1$, where N is the filter length. The impulse response of a discrete-time interpolation filter can therefore be obtained by sampling the underlying $h_a(t)$ at the desired instants, given by the fractional interval μ. Let us define $\mu_q(l)$ as the quantized fractional interval, that is, the fractional interval represented in finite precision format. The quantized fractional interval can be expressed as follows

$$\mu_q(l) = \frac{k_l}{2^B}\text{, with } 0 \le k_l \le 2^B-1, \tag{2}$$

where B is the number of bits. Thus, $\mu_q(l)$ has $K=2^B$ values and it is assumed that these quantization levels are uniformly spaced. Now the output of the hybrid analog/digital model is given by

$$y(l) = \sum_{k=-n/2}^{n/2-1} h_a((\mu_q(l)+k)T_{in}) \cdot x(n(l)-k) \tag{3}$$

Therefore, due to quantization, the filter response is replaced by the quantized $h_a((\hat{\mu}_l+k)T_{in})$. An example of such an impulse response for a polynomial-based interpolation filter is shown in Fig. 2 for $B = 3$.

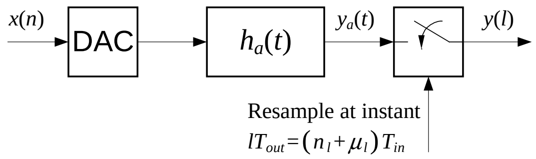

Fig.1. The hybrid analog/digital model for the interpolation filter.

As can be seen from Fig. 2, the impulse response of the quantized filter is obtained first by sampling the impulse response of the unquantized filter with the sampling interval of T_{in}/K and then by using the zero-order hold (ZOH) to reconstruct the continuous-time response. The effect of quantization in the frequency domain is shown in Fig. 3 for $B = 3$. The frequency response of the sampled filter has images at the frequencies KF_{in} and its multiplies. When these images are attenuated by the ZOH, we end up with the frequency response of the quantized interpolation filter.

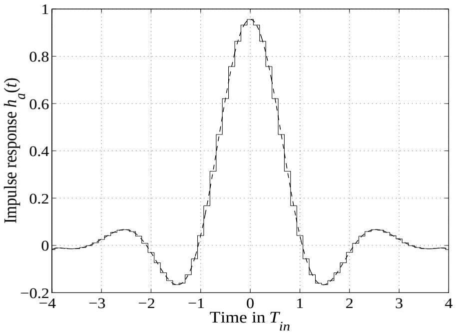

Fig. 2. Impulse response of the interpolation filter for unquantized (dashed line) and quantized (solid line) fractional interval, with B = 3.

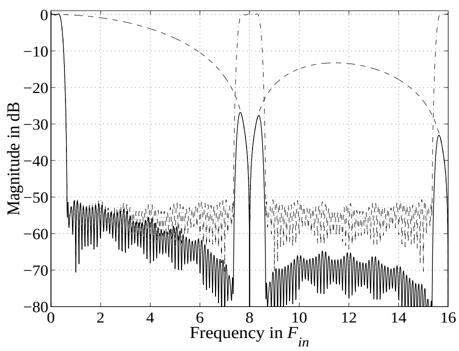

Fig. 3. The frequency response of the example interpolation filter when μl is quantized to B=3 bits.

The quantization of the fractional interval has several effects in the frequency domain. First, the ZOH causes some distortion to the passband of the interpolation filter, which does not usually need any compensation in other filtering stages. The most significant effect is caused by the attenuated images in the vicinity of the frequencies nKF_{in} for $n = 1, 2, 3, \cdots$: These high frequency components, which are called as quantization images, may fold into the baseband when the reconstructed signal $y_a(t)$ is sampled to get the output samples $y(l)$. The level of these quantization images are mainly affected by the number of bits, i.e., by $K = 2^B$ and the passband edge of the interpolation filter. The first quantization image has the highest value. The maximum value of the quantization images denoted by $\hat{I}_{max}$ can be approximated by

$$\hat{I}_{max} \approx \mathrm{sinc}(1 - \frac{f_p}{KF_{in}}), \qquad (4)$$

where f_p is the passband edge in Hz. Starting from Eq. (4) we can determine approximately the influence of the number of bits B and the input bandwidth, denoted by the passband edge frequency f_p on the amplitude of the quantization images. Eq. (4) can be approximated as follows

$$\hat{I}_{max} \approx 20\log(f_p/F_{in}) - 6.02 \cdot B \qquad (5)$$

From Eq. (5), where $\hat{I}_{max}$ is the maximum amplitude (or minimum attenuation) required, we can obtain the minimum value of B needed in our design as

$$B \geq (\hat{I}_{max} - 20\log(f_p/F_{in})) / 6.02 \qquad (6)$$

The key parameters effecting on the amplitude of the quantization images are the input bandwidth and the number of bits, shown in Fig. 4. As expected from Eq. (5), there is a logarithmic relationship between the bandwidth and the image amplitude. It also shows that a narrowband system will have a better performance in terms of quantization image attenuation for the same number of bits. As expected from Eq. (5), increasing the value of B by one bit, the attenuation is increased by about 6 dB. Thus, it would be easy to attenuate the aliased components as much as we desire by increasing B, but this parameter must be limited to reduce hardware costs.

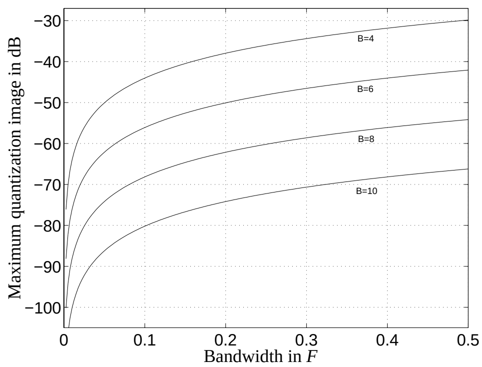

Fig. 4. Relationship between the input bandwidth and the attenuation of the first quantization image.

3. FRACTIONAL INTERPOLATION CASE

In this section, we will study the impact of the presence of interference signals on fractional interpolation filters when quantizing the fractional interval. Figures 5(a) and (b) illustrate the scenario when anti-imaging filters are used, which is a common choice in the design of sampling rate conversion to a higher output frequency. The disturbing signals are shadowed, and the frequency response of fractional interpolation filter is indicated with dashed line. Figures 5(c) and (d) show the case when an anti-aliasing filter is utilised. This latter filter sets harder requirements in terms of a narrower transition band, and therefore their implementation is more complex [Hen00].

In Fig. 5(b), the interference signals can fall into the quantization images, and consequently, in order to guarantee a proper attenuation of these undesired components, the amplitude of the quantization images must match the specifications of the stopband attenuation. In other words, the amplitude of the possible disturbing signals has an impact on the choice of the number of quantization bits for the fractional interval. This is because the quantization images are generated in the replicas of the passband and transition band regions, and in anti-imaging filters there might be interferences in the transition band.

However, when an anti-aliasing filter is designed, as depicted in Fig. 5(c) and 5(d), the possible disturbing signals as well as their images fall into the stopband region, and they are not affected by the appearance of the quantization images. Consequently, regardless of the amplitude of the interference signals, the attenuation of the quantization images is only related to the required attenuation of the replicas of the desired channel. Hence, the interference signals do not have any effect when dimensioning the wordlength of the fractional interval. Although this is an advantage of the anti-aliasing filters, anti-imaging filters are usually a better option due to their lower complexity when performing an increasing of the sampling rate. In the next section, we will study how the interference signals affect when fractional decimation is carried out using different implementations.

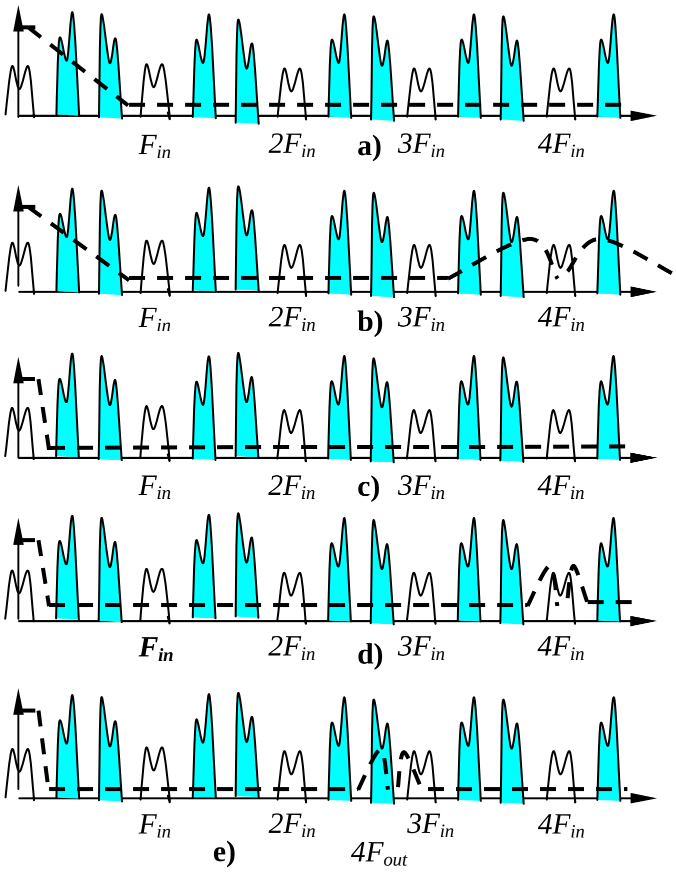

Fig.5. SRC with anti-imaging filter with (a) high precision and (b) finite precision of 2 bit; SRC with anti-aliasing filter with (c) high precision, (d) finite precision of 2 bits, and. (e) when μ is defined in terms of T_{out}.

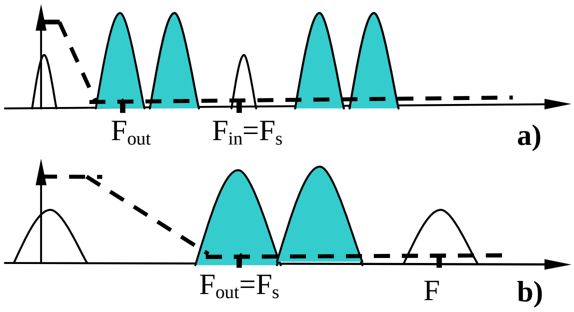

Fig. 6. Input signal spectrum and the anti-aliasing filter defined in terms of the (a) input and (b) output sampling interval.

4. FRACTIONAL DECIMATION CASE

When performing an increasing of the sample rate, both anti-aliasing and anti-imaging approaches can be used. In the previous section, effects of the quantization of μ in both types of filters were analysed. In this section we will concentrate on the reduction of the sample rate, or fractional decimation. When the system must perform a reduction of the sampling rate, the blocking signals must be attenuated because they can be aliased to baseband. Hence, anti-aliasing filters are needed.

As shown in Figure 6(a), if we intend to design such a filter as an interpolation filter, for instance by means of the Farrow structure, the transition band becomes very narrow, and consequently the filter complexity is increased. A better solution for this problem can be obtained if we can specify the problem from the output frequency point of view. In this approach, as long as the sampling rate is lower, the transition band gets wider, and consequently, the filter requirements are relaxed (see Figure 6(b), where the slope is less steep), leading to a reduction in the complexity of the system. In [Hen00] and [Bab01], two different implementations of a new approach to generate efficiently anti-aliasing interpolation filters by means of the so-called transposed Farrow structure are described. Here we will analyse the quantization effects in this kind of filters.

When the fractional interval is defined as a function of the output sampling interval T_{out}, which is the case of the trasposed Farrow structure, the quantization images are now centered at multiples of the output sampling frequency F_{out}, instead of multiples of F_{in}. This is illustrated in Figure 5 (e), where a conversion factor of 0.71 has been considered. In this case, an interference signal may fall into these quantization images, and it will be aliased to baseband. Therefore, the amplitude response provided by these quantization images must guarantee an acceptable attenuation for such disturbing signals. As long as the level of an interference signal can be much higher than the desired signal, this could establish much higher requirements for the attenuation of the quantization images, which in other words means that more bits are necessary to represent the fractional interval. In conclusion, when the fractional interval is referred to the output time interval T_{out}, the presence of strong blocking signals at the input of the interpolator has an impact on the final wordlength of the fractional interval.

5. BASIC EXAMPLE: MODIFIED AND TRASPOSED FARROW STRUCTURE

This example illustrates the effects of quantization when the interpolation filter is used for reducing the sampling rate by a factor of 2.49. The input signal is composed of a desired sinusoidal signal at $0.15F_{in}$ and a interference located at 0.35 F_{in}. This interference signal is 30 dB stronger than the desired baseband signal, and the specifications are set such that the maximum signal aliased to baseband must be 20 dB below the baseband signal. Thus, an attenuation of 50 dB is required. The anti-aliasing decimation filter is designed for both the modified Farrow structure, as an example of interpolation filter working at the input sampling rate, and transposed Farrow structure, which is a filter defined in terms of the output sample rate.

The filter specs are: passband edge at $0.15F_{in}$, stopband edge at $0.35F_{in}$, 50 dB of stopband attenuation and a passband ripple of 0.01. With these specs, the filter parameters for the modified Farrow structure are [Ves99]: filter length $N=14$ and the degree of the interpolation $M=4$.

Now we perform this decimation with the transposed structure. In this case, the filter specifications are given in terms of the output sampling rate, which is $1/\beta=2.49$ times lower. Thus, the passband edge is $0.3735F_{out}$, the stopband edge is located at $0.8715F_{out}$. With this specs, the designed filter has a length of $N=8$ and the degree of interpolator is M=3, which is a significant reduction of the complexity of the system, compared to the modified Farrow structure.

The designed filters meet both the desired specifications when μ is represented in high precision, with the desired frequency component, located at the frequency $0.15/\beta$, that is, 0.3725, and some aliased components adequately attenuated (below - 20 dB amplitude). The goal is to check the number of bits needed to represent μ so that the aliased components are below –20 dB. In theory, the filter designed with the modified Farrow structure is an anti-aliasing filter working at the input sampling rate, and the interferences do not affect the required attenuation of the quantization images, and 20 dB is enough to take care of the replicas of the desired signal. Figure 7(a) shows the output of the modified structure, with B=1 bit, which, as expected, is already good enough.

On the contrary, as explained in Fig. 5(e), in the trasposed Farrow structure the interferences can fall in the quantization images, so their amplitude must be below -50 dB. From Eq. (6) we can see that B=7 bits are needed in this case. Figures 7(b) and (c) illustrates the output of the trasposed Farrow structure when μ is quantized to B=1 and B=7 bits, respectively. By comparing Fig. 7(a) and (b) we can see the different effect of the interference in these two filters. Figure 7(c) confirms that 7 bits are sufficient.

6. CONCLUSIONS

In this paper, the quantization of the fractional interval for different type of filters has been addressed. It was shown that the requirements in terms of number of bits may vary considerably depending on the type of filter designed.

REFERENCES

[Bab01] D. Babic, T. Saramâki and M. Renfors, "Implementation of the transposed Farrow structure ", submitted to *Proc. Int. Conf. on Acoustics, Speech and Signal Proc., ICASSP2002*, Orlando, USA, May 2002.

[Hen00] M. Henker and G. Fettweis. "Continuous time digital filters for sample-rate conversion in reconfigurable radio terminals", *Proc. European Wireless*, pp.55-59, Dresden, Germany, Sept.2000.

[Laa96] T. I. Laakso, V. Välimaki, M. Karjalainen, and U.K. Laine, "Splitting the unit delay", *IEEE Signal Processing Magazine*, vol. 13, pp 30-60, Jan. 1996.

[Lop00] F. Lopez, J. Vesma, T. Saramaki, and M. Renfors, "The effects of quantizing the fractional interval in interpolation filters", in *Proc. IEEE NORSIG 2000*, Sweden, Jun. 2000.

[Ves99] J. Vesma, *Optimization and Applications of Polynomial-Based Interpolation Filters*; Doctoral Thesis, Tampere University of Technology, Publications 254, 1999.

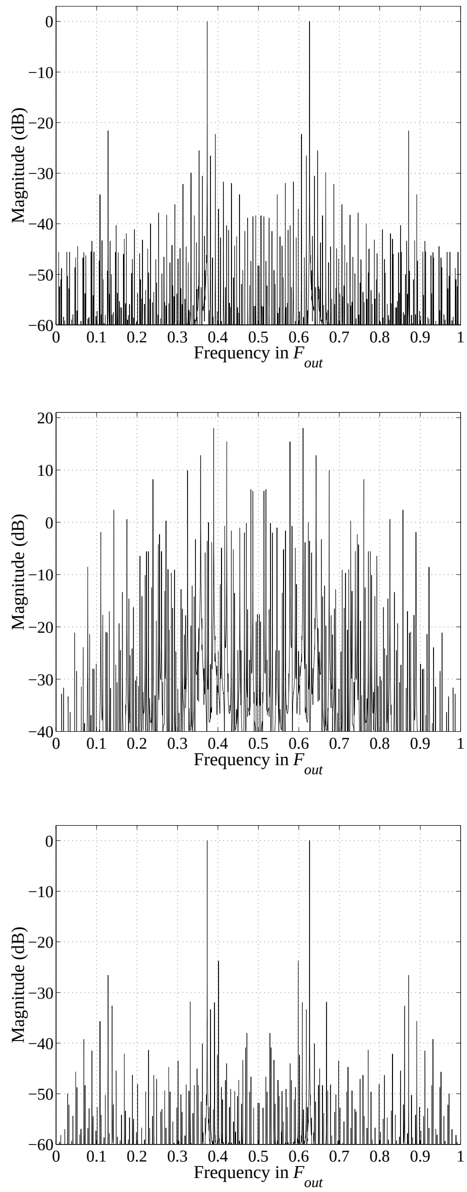

Fig. 7.The output spectrum for the modified Farrow structure for (a) B=1, trasposed Farrow structure for (b) B=1, and (c) B=7 bits.