

The Quest For The Perfect Resampler

2003.06.23

Laurent de Soras

Web: <http://ldesoras.free.fr>

ABSTRACT

This is a technical paper dealing with digital sound synthesis. Among all known sound synthesis techniques, there is at least one that will never be outdated: sample playback. We present in this document a computationally cheap method to play back samples at variable rate without perceptible aliasing. The main application of this technique is virtual music synthesisers running on personal computers.

KEYWORDS

Sound synthesis, sample playback, resampling, interpolation, alias free

0. Document structure

0.1 Table of content

0.	DOCUMENT STRUCTURE	2
0.1	TABLE OF CONTENT	2
0.2	REVISION HISTORY	2
0.3	GLOSSARY	3
1.	INTRODUCTION	4
2.	ROUND-UP OF CLASSIC RESAMPLING TECHNIQUES.....	5
2.1	KINDS OF INTERPOLATORS	5
2.2	MIP-MAPPING.....	6
2.3	OVERSAMPLED SAMPLE SOURCE.....	6
3.	PROPOSED SOLUTION	7
3.1	OVERVIEW	7
3.2	MIP MAPPING	8
3.3	INTERPOLATOR DESIGN	10
3.4	INTERPOLATING THE INTERPOLATOR.....	14
3.5	THE CASE $r < 1$	16
3.6	FINAL DOWNSAMPLING.....	18
4.	IMPLEMENTATION.....	19
4.1	A COMPLETE EXAMPLE.....	19
4.1.1	Requirements	19
4.1.2	MIP-mapping and data storage.....	19
4.1.3	MIP-map filter design.....	20
4.1.4	Interpolation filter design, $r > 1$	21
4.1.5	Design for $r < 1$	24
4.1.6	Downsampling filter.....	24
4.1.7	Transitions between MIP-map levels	25
4.1.8	Performances.....	25
4.2	POSSIBLE VARIATIONS	26
5.	REFERENCES	27
6.	APPENDICES	29
6.1	FILTER COEFFICIENTS	29

0.2 Revision history

Version	Date	Modifications
1.0	2003.06.20	Initial release
1.1		Miscellaneous typos

0.3 Glossary

CPU	Central Processing Unit
DAW	Digital Audio Workstation
FFT	Fast Fourier Transform
FIR	Finite Impulse Response.
IFFT	Inverse Fast Fourier Transform
IIR	Infinite Impulse Response
MAC	Multiply and Accumulate
MIP	Here, Multum In Parvo.
SDRAM	Synchronous Dynamic Random Access Memory
SIMD	Single Instruction, Multiple Data
SNR	Signal to Noise Ratio
TPDF	Triangular Probability Density Function

1. Introduction

The kernel of a sample-playback synthesiser (called sampler) is the resampling algorithm, using an interpolation filter. They have been already deeply studied. The main issues are the maximisation of the passband flatness and the optimisation of the SNR, which is degraded by the aliasing introduced by the interpolation. We will present in this document a method, or more exactly a set of methods, to interpolate sample data efficiently – in every sense of the term.

2. Round-up of classic resampling techniques

2.1 Kinds of interpolators

Interpolators generally fall into two classes: polynomial and FIR. The difference essentially lies in implementation, as polynomial interpolators can be generally described by their impulse response.

Low-order polynomial interpolators give surprisingly decent results at a low computational cost [1]. For example, the 4-point third-order Hermite interpolator does a very good job for a few multiplications and additions. One can represent the interpolation computation as a matrix multiplication. The 4,3 Hermite is expressed as:

$$f_{4,3H}(t+k) = [1 \quad t \quad t^2 \quad t^3] \times \frac{1}{2} \begin{bmatrix} 0 & 2 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ 2 & -5 & 4 & -1 \\ -1 & 3 & -3 & 1 \end{bmatrix} \times \begin{bmatrix} x_{k-1} \\ x_k \\ x_{k+1} \\ x_{k+2} \end{bmatrix} \quad (2.1-1)$$

To increase the quality, it is necessary to raise the order and the number of points. Unfortunately this leads quickly to large calculations, which are not compensated by a significant enough increase of quality.

On the other hand, we have the generic FIR convolutions. The first method is the direct application of the sampling theorem, where a window bounds the sinc function. This leads to multirate designs such as described in [4]. They have many pros and cons:

- Possibility to design the interpolation response curve accurately
- Straightforward technique working very well for all kind of resampling
- Only rates as fractions of integers (N/M) are handled.
- A lot of memory is wasted to store the complete impulse of a single N/M combination.
- Impulse should have a relatively important number of zero-crossings to get high quality (low aliasing, flat passband).
- Calculation cost increases proportionally with the resampling factor.

High quality interpolators are computationally cheaper than equivalent polynomial ones. This design is efficient for fixed rates, which is not the case in music synthesisers, where there is one rate per note, and it can change subtly if special effects such as vibrato or pitch bend are applied. A more advanced method is to linearly interpolate the impulse coefficient to save memory [5]. This solves the fixed and fractional rate problems, but the implementation is not as efficient as it could have been because of the linear interpolation on the impulse whose width may vary depending on the rate. This very specific point will be addressed later in this document.

We can give attributes to characterise response of polynomial and FIR interpolation:

- Transition bandwidth. The part in the passband affects its flatness. The part in the stopband creates aliasing.
- Amplitude of the first side-lobe. Lobes are images of the passband replicated in the higher zones of the spectrum, hence creating aliasing.
- Slope of the lobe top attenuation, can often be measured in dB/octave.

2.2 MIP-mapping


Additionally, interpolators are often used in conjunction with MIP-mapped samples. This technique is borrowed from computer graphics where it is used to reduce the aliasing when drawing textures on surfaces in 3D objects.

Document [8] gives a good definition: "MIP mapping features multiple images of a single texture map at different resolutions to represent surface textures at varying distances from the viewer's perspective: the largest scaled image is placed in the foreground and progressively smaller ones recede toward the background area. Each scale difference is defined as a MIP map level. MIP mapping helps avoid unwanted jagged edges (called jaggies) in an image that can result from using bit map images at different resolutions. MIP comes from the Latin *multum in parvo*, meaning a multitude in a small space."



MIP-mapped bitmaps.

Thus, the rendering engine has not to average the colours of a whole area of a distant texture to compute a given screen pixel. The right averaging is pre-calculated in each level. This is exactly the same principle with sounds. The original sample exists in many levels. Each is optimally low-pass filtered and decimated at a given sampling rate. Generally they are octave-spaced.



Five audio MIP-map levels, from 0 to 4.

MIP mapping reduces drastically the aliasing while upsampling with a high ratio, but high ratio become less and less needed with multi-layer sample banks, where each note of an instrument is sampled and stored separately. It remains entirely useful with wavetable synthesis where the "analogue" waveform shape is not affected by the pitch. However, MIP mapping does not completely remove aliasing, especially the one created by the interpolator.


2.3 Oversampled sample source

As depicted in [1] and [2], polynomial interpolator performances are dramatically improved on oversampled input data, where only a small part of the full spectrum is occupied. This gives a much flatter passband and very reduced lobes. The more oversampled, the better the performances. The compromise is an important increase of the room taken in memory. Oversampled source is often used partially along with MIP mapping, increasing SNR obtained with higher levels.

3. Proposed solution


3.1 Overview

The method we are going to describe is a three-stage process, consisting of MIP mapping, interpolation and downsampling. Suppose we want to resample by a factor r . Octave MIP mapping pre-filters and decimates the sample, allowing r to be high and avoiding expensive calculations. But if applied directly, an ideal interpolator would create aliasing by resampling by a local factor $r > 1$, or truncate the frequency content if $r < 1$.



Top figure shows the MIP map spectrums for levels 0 and 1. Below, the figures show both ideal and level-1 interpolations. Left: $r = 3$ ($r = 3/2$) and right: $r = 3/2$ ($r = 3/4$).

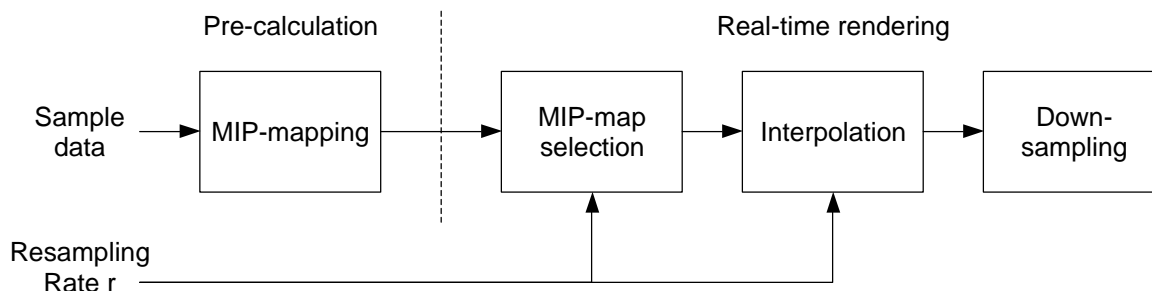
The interpolator runs at twice the sampling rate of the output, doubling the frequency content capacity. From now, we call r the ratio used to interpolate MIP map; it depends on the chosen level and the oversampling factor. Thus by choosing a MIP map level in order to make $r \leq 1$, Shannon principle is not violated. Aliasing can be theoretically avoided by low-pass filtering correctly during downsampling. And because final bandwidth is half the bandwidth of the 2x-oversampled signal, choosing $r \geq 1/2$ does not truncate the sample frequency content. Therefore the constraint is $r \in [1/2 ; 1]$. Spacing MIP maps by octave is sufficient to always find a MIP map and a r satisfying this constraint for a given resampling factor r .



Interpolation at 2x-oversampling. The dashed line shows the output Nyquist frequency.

The interpolator is a FIR. For maximum performances, it has a fixed bandwidth and is not scaled to lower the cutoff frequency in the case $r > 1$. Oversampling makes this kind trick useless, as we showed above. The FIR is stored in memory as a polyphase windowed impulse. Its coefficients are linearly interpolated in real-time to produce a very accurate interpolator.


The final stage is the 2x-downsampling. This is a polyphase implementation of an efficient elliptic IIR half-band filter and decimation, described in [9]. This is the only part of the process which is not linear-phase.



Resampler overview.

3.2 MIP mapping

MIP mapping is the first stage of the process. For the maximum efficiency, it has to be done off-line, once before any other processing. As we showed previously, optimal use of the system is reached when each MIP map level is spaced from its neighbours by an octave. The first level is obviously the original sample. To obtain every next level, we filter the current one with a half-band filter and decimate it by 2. The total amount of required memory does not exceed twice the original sample size.



Cascade of half-band filters and decimators producing MIP-map levels

It is important to keep the half-band filter linear-phase. Indeed, when switching from one MIP map level to another because of a variable resampling rate, a phase distortion between both levels would create a discontinuity in the signal. Thus, FIR filtering is the simplest method to proceed here. We need a good filter here because there is almost no headroom. We assume that calculation time is not important as the filtering is done once, during the system set-up. Filter has to be designed in order to make the passband as flat as possible, transition band the smallest and rejection band the lowest. Several methods can be used, such as Parks-McClellan algorithm [10] (also known as Remez Exchange) or sinc impulse truncated by a Dolph-Chebyshev window [7]. One or two hundred of taps would do the job correctly.

As filter data may be long, it is recommended to implement the convolution using a FFT with the overlap-add or overlap-save method [11]. Also, resolution issues must be taken into account. Indeed, storing the samples with the minimum necessary bitdepth and reusing them afterwards to generate the subsequent stages may accumulate roundoff errors. Instead, process every level simultaneously with maximum resolution, reducing the bitdepth only at the final stage to store the samples into memory. Here, dithering could help to save some bits of precision [12]. However, noise shaping is not recommended for the reasons given in [13].

Also, the FIR filtering introduces a delay. This delay has to be compensated either by truncating the beginning of each level or by re-indexing the sample arrays, in order to keep an uniform time reference between MIP map levels.

At playback time, the MIP map level index can be defined by:

$$l(r) = \max(0, \lfloor \log_2(r) \rfloor) \tag{3.2-1}$$

Where r is the resampling factor and $\lfloor x \rfloor$ denotes the greatest integer less than or equal to x – the floor function. 0 is the index of the original sample, 1 is 2x-downsampled, 2 is 4x-downsampled, etc. From that, we can deduce the local resampling factor \mathbf{r} :

$$\mathbf{r} = \frac{r}{q \cdot 2^{l(r)}} \tag{3.2-2}$$

Where q is the oversampling factor of the interpolation stage, $q = 2$ in our design. We can also deduce the number of levels to generate given the maximum value for r .

In the extreme case, in which the levels are produced down to reach a single-sample table, this MIP-mapping will occupy twice the original sample size in memory. Indeed:

$$S' = S + \frac{S}{2} + \frac{S}{4} + \frac{S}{8} + \dots = S \sum_{k=0}^{+\infty} \left(\frac{1}{2}\right)^k = S \frac{1}{1 - \frac{1}{2}} = 2S \quad (3.2-3)$$


In common use cases, and with the arrival of massive multi-sample instrument banks, algorithm does not need to reach such high sampling ratio. Often, only two levels are enough (0 and 1), occupying a much memory as one and a half times the size of the original sample data.

3.3 Interpolator design


We have chosen to implement the interpolator as a FIR because of its flexibility in response curve. SIMD instruction sets of the modern CPUs strongly improve the FIR performances by executing several multiplications and additions simultaneously, whereas they are less efficient on polynomial interpolator implementation.

Unlike many interpolators based on FIR design, ours has a fixed cutoff frequency. It is not really a low pass filter, because the cutoff is set to the Nyquist frequency of the sample (actually the current MIP map level). Indeed, we need to preserve the whole spectrum. Because of the oversampling, interpolated sample content will never go over the destination Nyquist frequency (F_{N1}). The goal here is to reconstruct the exact continuous signal and to sample it at the relative rate r .

Note that $r \leq 3/2$ is theoretically the real upper bound to avoid aliasing with 2x-oversampling. With this hypothesis, the maximum content frequency reaches virtually $3F_{N2}/2$ and aliases to $F_{N2}/2$. Therefore, the alias part of the oversampled signal remains always above $F_{N2}/2 = F_{N1}$ and will be completely discarded by subsequent downsampling. Or conversely, $3/2$ is a sufficient oversampling factor to use octave-spaced MIP-mapped samples. However we keep this headroom for the interpolator transition band, so it can extend to $3F_{N2}/2$ without audible effect. Moreover, $3/2$ x-oversampling requires a non-integer number of input samples per output sample, which can complicates things a bit if samples have to be outputted in odd-length blocks (it requires buffering). This problem does not appear with 2x-oversampling.



Ideal interpolation at $r = 3/2$. Aliasing stops above F_{N1} .




Real-life interpolator response at $r = 1$ and use of the headroom for the transition band.

So the FIR is basically assimilated to a variable fractional delay intended to reconstitute sample data for any fractional position. We put down:

- $p = \lfloor t \rfloor$ and $d = t - p$ the integer and fractional decomposition of a sample position $t = p + d$. Obviously, p is integer and $d \in [0;1[$,
- N an even integer representing the impulse response length - approximately the number of zero-crossings,
- $G_c(t)$ the interpolator impulse response, centred on $N/2$ and defined in the range $[0; N[$,
- $x[p]$ the sample data for the current MIP map level,
- and $x_c(t)$ the continuous function corresponding to the interpolated sample data.

We get:

$$x_c(p + d) = \sum_{k=1}^N x\left[p + k - \frac{N}{2}\right] \cdot G_c(N - k + d) \quad (3.3-1)$$



Sample interpolation for $\{p, d\} = \{2, 0.4\}$.
The grey curves show original sample's contributions to the sum.

$G_c(t)$ is a continuous function. Actually we can consider it as a set of discrete functions depending on a parameter d :

$$G_d[k] = G_c(k + d) \tag{3.3-2}$$


Therefore we use only one function $G_d[k]$ to compute the value of a position. The formula (3.3-1) becomes:

$$x_c(p + d) = \sum_{k=1}^N x\left[p + k - \frac{N}{2}\right] \cdot G_d[N - k] \tag{3.3-3}$$

The variable d is continuous. In concrete implementation, we cannot store in memory an infinite number of functions $G_d[k]$. For the practical implementation, we quantise d and assign it two basic parameters:

- The number of taps per phase. For a given fractional position, a phase of the FIR is selected and convolved with the sample data. This is the N parameter already mentioned.
- The number of phases M , equally spaced in the range $[0;1[$.


The N parameter directly impacts the steepness of the curve in the transition band and the ripples in the passband and stopband, as in every static FIR filter design. When all the phases are interleaved and sorted by ascending time, it gives a global impulse whose cutoff frequency is F_N/M (frequency related to the MIP-map level).



Top: all the $G_d[k]$ curves for $M = 4$. Below, the interleaved impulse.

So we are now focused on the design of a “brick-wall” low-pass filter of length $N \times M$. There are several methods to do it. The preferred way is probably the Parks-McClellan algorithm, but the simplest is probably the windowed sinc and works well for long FIR. The Dolph-Chebyshev window provides equiripple in the stopband and good transition bandwidth properties. Note that the stopband ripple specification provides guaranteed attenuation, but in concrete case, it may be a bit too much. One has to check the actual frequency response with an FFT and adjust recursively the stopband specification to optimise the transition bandwidth.

Another thing to consider: we left headroom by constraining the oversampling rate and r value range. This allows the main lobe to extend to $3F_{N2}/2$. So we can shift up the filter cutoff in order to let the main lobe reach this frequency. Thus, we maximise the flatness of the passband.



Taking advantage of the headroom to improve filter flatness in the passband.
Left figure: $r = 1/2$, right figure: $r = 1$.


One more improvement can be done, considering that the passband does not extend exactly up to Nyquist. Real-world requirements would include a little headroom between the end of the passband and Nyquist frequency. Indeed, the very top end of the spectrum is generally filled with garbage produced by non-brickwall anti-aliasing filters. We put F_{CP} as the passband end, as fraction of the cutoff frequency, here F_{N1} . Taking the aliasing into account for $r = 1$, we can calculate the optimal stopband starting frequency:

$$F_{CS} = \frac{4 - F_{CP}}{2} \tag{3.3-4}$$

With $F_{CP} = 0.9$, which is a rather common requirement, we obtain $F_{CS} = 1.55$. If Parks-McClellan algorithm is used, these frequencies bound the pass- and stop-bands.

3.4 Interpolating the interpolator

Obviously quantising the number of functions G is equivalent to quantise the fractional position d . As shown in [3], the number of phases (M) influences the SNR with a rate of 6 dB/bit. Below is the compared responses of such a quantised interpolator, with $N = 16$ and different values for M :



Impulse response of discrete and continuous interpolators

It clearly shows a huge gap between the continuous and discrete versions of the interpolator. The continuous interpolator's side-lobes are only caused by the impulse

windowing whereas origin of the main side-lobes of the discrete interpolators is the phase bit reduction. In these conditions, we have to increase M to unacceptably high values to get a decent SNR. That is why it is necessary to interpolate correctly the functions $G_d[k]$. The different phases constituting the whole impulse can be numbered from 0 to $M - 1$. We introduce the integer variable q :

$$q = \lfloor d.M \rfloor \quad (3.4-1)$$

The way we round $d.M$ is not the most appropriate to use directly $G_{q/M}[k]$ – it would have been better to round it to the nearest integer, halving the phase error and gaining 6 dB of SNR. Our choice is the linear interpolation, so this rounding is actually a logical move. So we finally define the interpolated function $G_{id}(t)$ as:

$$G_{id}[k] = G_{\frac{q}{M}}[k] + (d.M - q) \left(G_{\frac{q+1}{M}}[k] - G_{\frac{q}{M}}[k] \right) \quad (3.4-2)$$

Formula (3.4-2) makes use of $G_1[k]$, which is not defined. According to the (3.3-2) definition, we extend the range of the class of G functions to ensure the coherence of (3.4-2):

$$G_1[k] = G_c(k + 1) \quad (3.4-3)$$


Therefore:

$$G_1[k] = G_0[k + 1] \quad (3.4-4)$$

This class of functions $G_{id}[k]$ gives us the possibility to define a new continuous function for the interpolator:


$$G_{c2}(k + d) = G_{id}[k] \quad (3.4-5)$$

Below, $G_{c2}(t)$ for $N = 16$ and $M = 4$. We evaluate its frequency response and compare it to the quantised one. Obviously $M = 4$ is too low for a real-world application, but the purpose here is to show clearly the principle and its benefits.



Frequency responses of $G_{C_2}(t)$ and its discrete version for $N = 16$ and $M = 4$.
Shaded curves show the effect of aliased side-lobes.

With $M = 128$, the aliasing is now rejected in the -100 dB oblivion:



Frequency responses of $G_{C_2}(t)$ and its discrete version for $N = 16$ and $M = 128$.

This new interpolator scheme is equivalent to a polynomial interpolator in presence of oversampled input, where M is the oversampling factor. If you need extra-low aliasing and care more about memory than computational cost, linear interpolation may not be enough. See [1] and [2] for deeper discussion on this topic.

3.5 The case $r < 1$

We have essentially dealt with the aliasing problem caused by the case $r > 1$. We took advantage of the oversampling to lower the requirement on interpolation filter transition bandwidth. In conjunction with MIP-mapping, we could have achieved a very steep filter for a relatively low real-time cost.

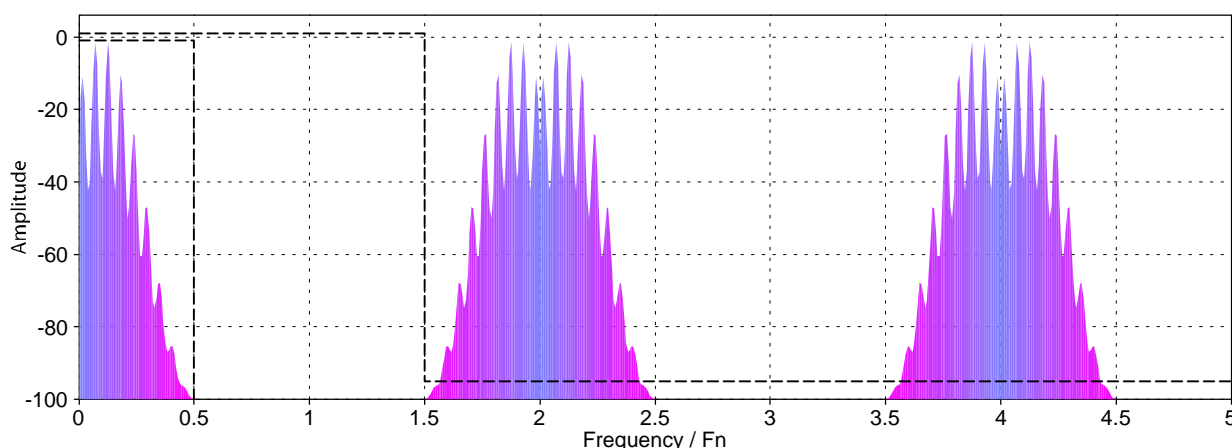
However things are a bit different when r is less than 1. Indeed, the interpolation filter slope is not as steep as it should be. The top end frequency content, which is naturally mirrored and replicated over $r.F_N$, is not filtered enough. Therefore we need to use a more steep filter. Actually, we have two solutions: post-filtering the signal with an efficient IIR (elliptic) low-pass filter, or make the interpolation filter stronger to reach a much thinner transition band.

The first option is very attractive. Elliptic IIR filters are quite cheap and we can obtain a steep slope for a relatively low number of MACs. But it has some drawbacks: lack of phase continuity with the case $r > 1$, problems with cutoff frequencies close to Nyquist and recursive structure preventing the parallelisation of operations. None of these problems are insoluble, but add significant complexity to the solution.

The second option may appear outrageously computation-greedy. Actually this is not totally true. With $r < 1$, we are not facing anymore the aliasing problem, occurring even using an ideal low-pass filter for interpolation. Here, aliasing would have been caused only by a non-brickwall filter, letting pass too much mirrored frequencies, some of them exceeding Nyquist limit. If the filter slope is steep enough, we do not need the oversampling trick, thus saving half the calculation time. This time can be reused to achieve this very brickwall filter we need.


If this solution is simple and quite efficient, it is far from perfect. Indeed, the brickwall-ness of our filter has its own limits. With twice as many taps, we can only reduce the transition band by a factor of two, not much more. This solution is acceptable if we have quite large transition bandwidth requirement, or if we are ready to make a compromise against ripples and attenuation.

However we have another asset, one more compromise to make. This one is on memory requirement. Our interpolation filter specifications are relative to the original signal, not directly its frequency content. If the latter does not occupy the full bandwidth, spectrum of the interpolated result will have "holes", corresponding to the absence of content in the original signal. By pre-oversampling sample data, we can reduce the areas where the filter has to actually filter, thus producing better interpolators. This is the topic of [1] and [2] regarding polynomial interpolation, but we can apply the same theory to FIR interpolators.



Interpolation filter requirement for 2x-pre-oversampled data.

How much do we need to oversample the input? Not much. Every additional sampling bandwidth immediately benefits to the transition band. Therefore 2x-oversampling is largely enough. If memory is a crucial concern, we can reduce the oversampling factor down to about $4/3$. Indeed the filter designed for $r > 1$ has roughly a $F_N/2$ transition bandwidth. This is the same width as the free band left by the mirroring of the spectrum after interpolation, between $3F_N/4$ and $5F_N/4$.



Interpolation filter requirement for 4/3x-pre-oversampled data.

Thus, we add one more oversampled MIP-map level. We could be tempted to use oversampled MIP-map everywhere, especially for the case $r > 1$. But doing this, we need to increase r in order to raise the significant band. This means that real-time oversampling allows less headroom for the filter, so more taps, more computations are required. Exactly the opposite of what we want...

We can also mention a third possibility: not changing anything to the algorithm. It is possible to live with the mirrored spectrum top, and it gives the possibility to get rid of the oversampling as soon as $r < 2/3$.

3.6 Final downsampling

The last thing we have to do is to get back from 2x-oversampling to the final sampling rate. We will use the polyphase IIR half-band filter described in [9] because of its efficiency. The phase distortion is less important here, and anyone needing to keep it linear can use classic polyphase FIR filters; this kind has been already described in the literature [4].

Theoretically, we have not any more headroom for the transition band here. However, sound signal is never critically sampled. A small part of the spectrum cannot be heard and the band above 20 kHz up to Nyquist frequency may be altered without noticeable effect. For example, sampling rate of 44.1 kHz leaves us a 4.1 kHz range for the transition between passband and stopband. Note that this headroom can also be used to widen a bit the transition band of the MIP-map and interpolation filters.

4. Implementation

4.1 A complete example

4.1.1 Requirements

In this section, we will design an interpolator built with the following specifications regarding audio quality:

- Passband defined from 0 to $\frac{9}{10} F_{N1}$
- Aliasing attenuation below -85 dB in the passband
- Passband ripple below 0.1 dB

We can also add obviously these loose constraints:

- Low memory use
- Low CPU use
- Real-time rendering
- Time-varying resampling rate

At 44.1 kHz sampling frequency, the passband extends up to 19.8 kHz, which is more than enough for audio-range applications. We will not give a magic formula to produce the best interpolator for a given specification. Instead, we can proceed by trial and error. It works fine within the frame of the method described above. The result will be probably not optimal, but will have a very simple and efficient implementation, hard to outclass given the specification.


4.1.2 MIP-mapping and data storage

As suggested in the previously described method, we will work with octave-spaced MIP-mapped samples and $2x$ -oversampling.

The MIP-map data type depends on the input data. In case of integer 16-bit samples, there is no real need to go to 32-bit integer or floating point because the noise floor is already set to approximately -96 dB (16-bit resolution). However one needs to be careful. If original data fits exactly in the data range, reaching 0 dB level, the resampled data will likely overflow. Take the worst case: the signal is a constant -1.0 value, with two consecutive samples jumping down at -1.0 , then going back to 1.0 . A perfect interpolation would produce important ripples on the positive side (on the negative side too, but lesser), reaching a maximum value:

$$\begin{aligned}x_{\max} &= 2 \left(\operatorname{sinc}\left(\frac{1}{2}\right) + \operatorname{sinc}\left(-\frac{1}{2}\right) \right) - 1 && (4.1.2-1) \\ &= 4 \operatorname{sinc}(1/2) - 1 \\ &\approx 1.5465\end{aligned}$$

So it is recommended to keep at least 3.8 dB of headroom regarding this issue, or more simply, augmenting the data capacity by one bit. In our case, because final SNR will be 85 dB, it is possible to use 16-bit integer MIP-map data, based on original 16-bit normalised sample data, reduced to 15-bit resolution in order to gain the necessary 1-bit headroom.



The need for headroom in signal interpolation.

4.1.3 MIP-map filter design


We use the Parks-McClellan algorithm to design the half-band filter ($F_C = F_N/2$) in order to match the following prototype for the frequency response:

- 1 between 0 and $9F_C/10$, the passband.
- 0 between $11F_C/10$ and F_N , the stopband.
- Weight of 100 for the stopband ripples, relative to the passband ones.

We obtain the following filter characteristics:

- Number of coefficients: 81
- Width of the symmetric transition band: $F_N/10$
- Passband ripple: 0.04 dB
- Stopband attenuation: -92.2 dB
- Overall group delay: 40 samples (linear phase)

The ripples are lower than the requirements because we have to take into account the subsequent passband pollution by further processing.



MIP-map filter frequency response

To filter the sample, we will use FFT to do the convolution in order to achieve smallest set-up time as possible. We must use a FFT length greater than 96. With 128, we can process only 32 samples per pass, which may be a waste of performance. 256-point FFT seems more appropriate. Also, 192 points are fine for any radix-3 FFT algorithm. The filter can be stored in the frequency domain, ready to be multiplied with frequency data, to avoid unnecessary operation. Again, we need to add a word about resolution. 16-bit integer FFT will not be enough, it is necessary to compute everything with at least 24-bit integer or 32-bit floating-point data and convert back to 16 bits at the final stage. Bit reduction will be done with a simple TPDF dithering.

4.1.4 Interpolation filter design, $r > 1$


We now need to define the most important parameters, N and M . N influences the overall quality of the interpolator, and more exactly the transition bandwidth. M regulates the level of aliased images caused by the interpolation of the interpolator. Given our experiments, a combination of $N = 12$ and $M = 64$ are well suited to match the requirements.

Again, we use the Parks-McClellan algorithm to design the interpolation filter, whose cutoff is located at $F_c = F_N / M$. We deduce the filter specification from equation (3.3-4) and feed the algorithm with the following prototype:


- Number of coefficients: $N.M - 1 = 767$
- 1 between 0 and $9F_c / 10$, the passband.
- 0 between $31F_c / 20$ and F_N , the stopband.
- Weight between passband and stopband determined by -0.08 dB / -85 dB ripples.

We specified only 767 coefficients to the algorithm. This is motivated by a phase issue: it is easier to operate with an odd-order filter, because the group delay is integer, thus easily controlled. It is possible to achieve the same results with an even-order filter, but it needs additional phase tweaking during the interpolation to smooth transitions between the MIP-map levels. The resulting filter has these properties:

- Transition band between: $9F_c / 10$ and $31F_c / 20$
- Passband ripple: 0.08 dB
- Stopband attenuation: -85.0 dB
- Overall group delay: 384 samples (linear phase), 6 with individual $G_d[k]$ functions.



Frequency response obtained using linear interpolation.

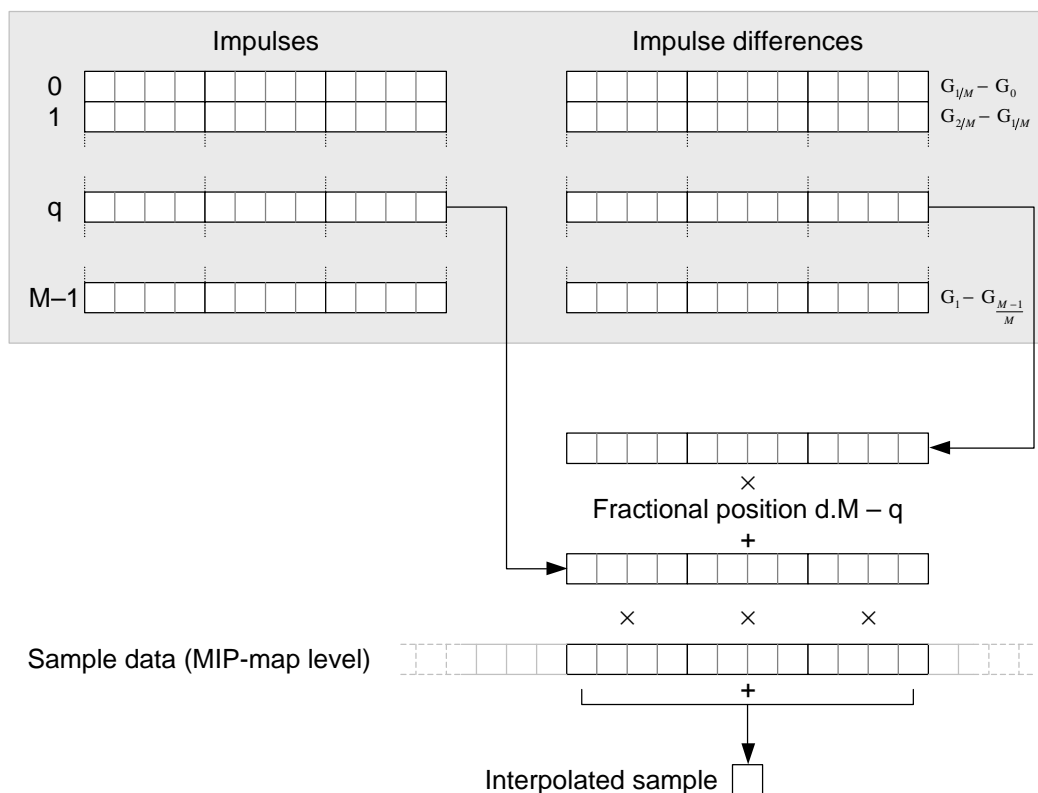


Zoom on the passband.


We now need to deal with the storage of the impulse in memory. The most efficient method is probably to store it as individual, de-interleaved $G_d[k]$ functions. Each one represents a FIR impulse, which can be read from memory and block-wise computed with MIP-map data, using SIMD processor instructions. Because we need to perform a linear interpolation – equation (3.4-2), it is smart to pre-calculate the difference between two consecutive $G_d[k]$ functions. The last difference function, corresponding to $G_1[k]$, is obtained by shifting $G_0[k]$ one sample to the left as shown in (3.4-4), and completing with a 0 at the end.

We store the tables by interleaving the Difference functions with the Impulse functions in order to take advantage of the SDRAM “burst mode”. Indeed, most DAWs run samplers with other effects or virtual instruments. Each virtual device of the audio network outputs a N-sample block when comes it turns. Thus data may not be present in the cache at the beginning of the processing, generating a cache miss and requiring a SDRAM access. Grouping simultaneously accessed data helps to minimize the slowing down.

Regarding the data format, impulse can be stored as 16-bit integers if normalized properly before. Our tests have shown that it has a minor impact on stopband rejection, diminishing it by a half dB, which is acceptable here. A bit-flipping method based on genetic algorithms such as Differential Evolution [14] may help to fix the problem. Actually the choice between 16-bit and 24- or 32-bit data is more a matter of implementation and performance, specific to the architecture, rather than a quality issue. However 16-bit accumulators are not enough for the filtering, they would require a few bits more to guarantee 85 dB of SNR.



Interpolation. Grey-background rectangle represents the program memory, containing the impulses $G_d[k]$ and their differences. Bottom part of the diagram shows the linear interpolation of the FIR, followed by its convolution with the sample. Data are grouped by packets of 4 samples, to show the possible use of SIMD instructions and the processing parallelisation. Note that one has to revert impulses and differences in order to index them correctly (the $N/2 - k$ of the convolution formula).



Effects of 16-bit coefficient quantification on interpolator frequency response.

4.1.5 Design for $r < 1$

As we stated in our analysis, the case $r < 1$ has a pretty efficient solution: using an additional MIP-map level, oversampled at a rate $4/3$. This is probably the ideal solution for a system where memory use and bandwidth are not an issue. However this is not always the case. Switching quickly between $r < 1$ and $r > 1$, i.e. in a vibrato, or playing a chord will consume a lot of memory bandwidth. We can reduce it by not adding this MIP-map level. Moreover the implementation adds a bit of complexity. Mainly for the sake of simplicity, we are going to look at the other solution we pointed: a steeper filter.

To keep the computational cost well balanced between both resampling ratios, we choose a double filter length without oversampling. Here we are cheating with the original interpolator characteristics we specified. Indeed the steepness is not enough and we cannot guarantee a perfect 90 % bandwidth. We need to trim it a bit, and to make the attenuation more progressive. The Parks-McClellan algorithm is set up with the following specifications for $N' = 2N = 24$ and $M = 64$:

Range	Weight	Actual ripple or attenuation (dB)
$[0 ; 0.85F_C]$	1	0.1
$[1.14F_C ; 1.99F_C]$	64	-81
$[2F_C ; 3.99F_C]$	128	-87
$[4F_C ; F_N]$	256	-93

4.1.6 Downsampling filter

The half-band IIR filter is designed with these specifications:

- Stopband attenuation: -90 dB
- Passband width: $9F_C/10$

It results in these effective characteristics:

- 7 coefficients
- Width of the symmetric transition band: $F_N/10$
- Effective stopband attenuation: -93.3 dB
- Insignificant passband ripple
- Group delay after decimation:
 - 18 samples at $9F_{N1}/10$,
 - 2.3 samples at $F_{N1}/2$,
 - 1.7 samples close to 0 Hz.

The group delay looks impressive in the spectrum top end, but actually attenuates quickly as frequency decreases. The coefficients support very well the quantification, therefore it is acceptable to store them in 16 bit if required. Anyway it is probably better to keep a higher resolution for state variables.

For the case $r < 1$, we do not oversample, but need to ensure the phase continuity with the other case. We proceed by feeding the downsampling filter with the non-oversampled signal, in which we insert zeros between each two samples.

4.1.7 Transitions between MIP-map levels

Changing the resampling ratio r sometimes requires changing the MIP-map level currently in use. Theoretically, the change is almost transparent in the final signal; both frequency contents are the same at the frontier. However, it makes harmonics appear or disappear instantaneously in the oversampled part of the spectrum, resulting in discontinuities. These discontinuities are broadband and impact the unfiltered, low part of the spectrum, producing audible clicks.

We have two easy solutions to workaroud the problem. The first one consists in not changing the MIP-map level once it has been selected, assuming the pitch changes are slight, and the aliasing or high frequency loss being considered as negligible.

Second solution involves cross-fading signals produced by both MIP-map levels. This method actually uses the first solution (MIP-map used out of its own pitch range) with the previous and next MIP-map. The time needed to run a full cross-fade requires limiting the rate of the pitch changes, or to delay subsequent level changes after the end of the current cross-fade. The major drawback of the cross-fading method is the doubling of the calculations.

Note that transition between $r < 1$ and $r > 1$ is exactly the same issue as filter characteristics change.

4.1.8 Performances

Regarding pure real-time DSP performance, the count per output sample is easy. Interpolator FIR is made of 12 taps, therefore 12 MAC. To generate this interpolator, we need to interpolate two impulses with the formula $y = x_1 + k(x_2 - x_1)$. Because $x_2 - x_1$ is pre-calculated, it involves one MAC per tap. So it is 12 MAC again, 24 MAC total. The interpolation filter is running at 2x-oversampling, we need to multiply this by 2. Then the downsampling filter requires 1 MAC per coefficient and output sample, so it is 7 MAC here.

Grand total is 55 MAC per output sample.

Obviously this cannot measure the real performances of the algorithm. We need to take into account the logic: MIP-map level selection, phase indexing, data loading and storing, etc. Anyway they are quite light in comparison with the calculation. Because of the vector structure of the data, dependencies between intermediate results are relatively low. Therefore it is generally possible to operate logic and computation in parallel. Moreover, it is the perfect opportunity to use SIMD instruction sets.

If many voices have to be mixed together, it may be more efficient to mix them before downsampling, although it is counterintuitive.

Besides the MIP map load, memory occupation can be roughly expressed by the following formula:

$$S = 2(N + N').M.D \tag{4.1.8-1}$$

Where D is the coefficient size, in byte; i.e. 4 for 32-bit floating point data. In our case, interpolator coefficients occupy 18 kB of memory.

We successfully implemented the proposed method and came out with an average speed of 178 clocks per sample on a classic desktop computer, which is largely enough for real-time sample playback. This was done in realistic testing conditions, taking cache issues into account. The program was written in pure standard C++, full 32-bit floating point processing,

without any assembly code or SIMD instruction. Thus, by optimising the critical parts using these programming techniques, one can effortlessly triple or quadruple the data throughput.

4.2 Possible variations

In the previous section, we described an implementation of our method. This is not the only possible one. We can tweak several parameters to fit particular requirements. In this section we will list – not exhaustively – some possible alterations:

- **Number of phases (M).** Increasing it rejects the aliasing caused by interpolation side-lobes, but requires more memory for the filter.
- **Phase length (N).** Increases the overall filter quality, the compromise between ripples/rejection and width of the transition band.
- **Phase interpolation.** We choose a linear one, but one also can use higher order polynomials to save memory against calculations.
- **Filter design method.** It depends on the tools you have got! If you cannot use the Parks-McClellan algorithm, go for windowed sinc. Classic fixed windows such as Blackman-Harris have generally too high side-lobes; prefer parametric windows if possible, like Dolph-Chebyshev or Kaiser.
- **Oversampling ratio.** Reducing it may reduce the computation, but requires steeper interpolation filters. In the other hand one can oversample at a higher rate to lower the filter requirements.
- **MIP-map level spacing.** Octave is easy to implement, but it is possible to choose double- or half-octave spacing (or other ratio), changing the memory requirement. The width of the filter transition band may be reduced, as well as the oversampling ratio.
- **Filter storage.** Because of the symmetry of the impulse, one can store only half of it in memory. However it is more complicated to fetch it at run-time.
- **Phase pre-interpolation.** Similarly, we could refuse to store the difference between two phases, to save calculations or memory.
- **Downsampling filter.** It can also be tweaked if necessary. Especially if the oversampling ratio is not a power of two, it may be useful to use a FIR instead of an IIR. FIR filters are particularly efficient for low resampling ratios.
- **Case $r < 1$.** Alternative methods have been evocated in a previous section.
- **MIP-map switching.** Same as above.
- **Output sample rate.** This is an important point as filter requirements can be completely different depending on the output sample rate. 44.1 kHz leaves 2 kHz of headroom, whereas 48 kHz doubles this headroom. Not speaking about extended rates such as 88.2 kHz, 96 kHz or higher. Whether to widen the transition band is a matter of requirements on the passband.

5. References

- [1] **Polynomial Interpolators for High-Quality Resampling of Oversampled Audio**
Olli Niemitalo, August 2001
<http://www.student.oulu.fi/~oniemita/DSP/deip.pdf>
- [2] **Performance of Low-Order Polynomial Interpolators in the Presence of Oversampled Input**
Duane K. Wise & Robert Bristow-Johnson
AES 107th convention, September 1997
- [3] **The Effects Of Quantizing The Fractional Interval In Interpolation Filters**
Jussi Vesma, Francisco Lopez, Tapio Saramäki & Markku Renfors
http://www.es.isy.liu.se/norsig2000/publ/page215_id108.pdf
- [4] **Multirate Digital Signal Processing**
Ronald E. Crochiere and Lawrence R. Rabiner
Prentice-Hall
- [5] **Digital Audio Resampling Home Page**
Julius O. Smith III, January 2000
<http://ccrma-www.stanford.edu/~jos/resample/resample.html>
- [6] **Windowing Functions Improve FFT Results**
Richard Lyons
Test & Measurement World, June/September 1998, pp. 37-44
<http://www.e-insite.net/index.asp?layout=article&articleId=CA187572>
<http://www.e-insite.net/index.asp?layout=article&articleId=CA187573>
- [7] **The Dolph-Chebyshev Window – A Simple Optimal Filter**
Peter Lynch, January 1996
<http://www.maths.tcd.ie/~plynch/Publications/Dolph.pdf>
- [8] **MIP-mapping**
<http://whatis.techtarget.com>
- [9] **Polyphase Filter Designer in Java**
Artur Krukowski
<http://www.cmsa.wmin.ac.uk/~artur/Poly.html>
- [10] **FIR Digital Filter Design Techniques Using Weighted Chebyshev Approximations**
L. R. Rabiner, J. H. McClellan, and T. W. Parks
Proc. IEEE 63, 1975
- [11] **Computational Improvements To Linear Convolution With Multirate Filtering Methods**
Jason R. VandeKieft, April 1998
<http://www.music.miami.edu/programs/mue/Research/jvandekieft>
- [12] **The Mathematical Theory of Dithered Quantization**
Robert A. Wannamaker, Ph.D. thesis, July 1997
Department of Applied Mathematics, University of Waterloo, Canada
<http://audiolab.uwaterloo.ca/~rob/phd.html>

[13] Whither Dither - Experience with High-Order Dithering Algorithms In The Studio

James A. Moorer and Julia C. Wen
AES 95th convention, October 1993

<http://www.jamminpower.com/main/articles.jsp>

[14] Differential Evolution homepage

Kenneth Price and Rainer Storn

<http://www.icsi.berkeley.edu/~storn/code.html>

