



HAL
open science

Reliable verification of digital implemented filters against frequency specifications

Anastasia Volkova, Christoph Lauter, Thibault Hilaire

► **To cite this version:**

Anastasia Volkova, Christoph Lauter, Thibault Hilaire. Reliable verification of digital implemented filters against frequency specifications. 2017. hal-01432000v1

HAL Id: hal-01432000

<https://hal.sorbonne-universite.fr/hal-01432000v1>

Preprint submitted on 11 Jan 2017 (v1), last revised 24 Apr 2017 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Reliable verification of digital implemented filters against frequency specifications

Anastasia Volkova, Christoph Lauter, Thibault Hilaire
Sorbonne Universités, UPMC Univ Paris 06, UMR 7606, LIP6
4, place Jussieu 75005 Paris, France
Email: first_name.last_name@lip6.fr

Abstract—Reliable implementation of digital filters in finite-precision is based on accurate error analysis. However, a small error in the time domain does not guarantee that the implemented filter verifies the initial band specifications in the frequency domain. We propose a novel certified algorithm for the verification of a filter’s transfer function, or of an existing finite-precision implementation. We show that this problem boils down to the verification of bounds on a rational function, and further to the positivity of a polynomial. Our algorithm has reasonable runtime efficiency to be used as a criterion in large design space explorations. We ensure that there are no false positives but false negative answers may occur. For negative we give a tight bound on the margin of acceptable specifications. Numerical results illustrate the application of our approach to comparison of transfer function design tools and of the various finite-precision implementations.

I. INTRODUCTION

The great majority of signal processing or control algorithms are designed on our desktop computer, and then embedded in digital devices, such as general purposes processors, DSPs, FPGAs or ASICs.

Linear filters are basic bricks of such algorithms, and they are usually designed using software like Matlab¹ or SciPy² which are used to compute filters fulfilling a given frequency specification. These tools rely on algorithms based on double precision Floating-Point arithmetic. Then, the filter is implemented, often with Fixed-Point arithmetic, for cost or power consumption reasons. The filter algorithm used for the implementation may directly round and use the coefficients obtained at design stage, or deduce new coefficients to round and use, from the previous stage. In both cases, the finite precision used in design and implementation stages induces some uncontrolled errors that may make the implemented filter/controllers differ from the initial specification. Currently, engineers have no other choice than to trust these software tools and perform some simulations to somehow test that the implemented filter respects its specifications.

So this article deals with the reliable *a posteriori* verification of the filter implementation against its mathematical frequency specification.

This problem can be decomposed in two parts:

¹<http://www.mathworks.com/>

²SciPy is a Python-based open-source scientific software for mathematics, science, and engineering: <https://www.scipy.org/>

First, suppose, we have a transfer function with its coefficients expressed as fixed- or Floating-Point numbers at some precision (integers scaled by powers of 2) and a set of band specifications. Then, we need to verify whether this filter satisfies the specifications. Such a problem boils down to verification whether a rational function is between bounds. We propose an algorithm that guarantees that no false positive answer is ever given, though some extremely rare false negatives are possible.

Secondly, suppose, we have an algorithm realizing the filter, defined a set of equations or as a data-flow graph (like those in Matlab/Simulink). In general case, the algorithm’s coefficients have been quantized and are hence not those of the transfer function. Therefore, to verify whether the filter satisfies band specifications, its transfer function must be computed. Usually, this cannot be done exactly. We propose a rigorous approach which computes a multiple precision approximation on the filter’s transfer function with a reliable error-bound. Then, we can check whether the implemented filter satisfies band specifications while taking into account the error of the transfer function computation. Moreover, if the filter does not satisfy the band specifications, our algorithm computes a tight extension to the frequency bands that the filter does satisfy.

The paper is organized as follows. Section II reminds some propositions about linear filters and defines the band specifications. In Section III, the reliable verification of a transfer function is exhibited. With the multiple precision approximation on the transfer function of any linear algorithm, Section IV checks whether an implemented filter satisfies its initial specification. Some examples illustrate the method are given in Section V before conclusion.

Notation: throughout the article matrices are in uppercase boldface, vectors are in lowercase boldface, scalars are in lowercase. As usual in the digital filter community, the complex unit is notated j ($j^2 = -1$). The conjugation operator on a complex number z is notated z^* .

II. PRE-REQUISITES

A. Digital filters and LTI systems

Digital filters are computational blocks that transform a signal u (i.e. a sequence $u(k)$, where $k \in \mathbb{Z}$ stands for the step-time, according to the sampling period) into a signal y . In this article, we focus on discrete-time Linear Time Invariant (LTI) filters only, i.e. filters that are linear and for which a time

shift (a delay) of the input sequence causes a corresponding shift in the output sequence [1], [2]. They are standard basic bricks of digital filtering (and control) and include FIR (Finite Impulse Response) and IIR (Infinite Impulse Response) filters.

The output of a LTI filter can be expressed by

$$y(k) = \sum_{i=0}^n b_i u(k-i) - \sum_{i=1}^n a_i y(k-i), \quad \forall k \in \mathbb{Z}, \quad (1)$$

where $\{a_i\}_{1 \leq i \leq n}$ and $\{b_i\}_{0 \leq i \leq n}$ are real numbers defining the filter (n is the order of the filter).

Instead of an input-output relationship in time (the output $y(k)$ with respect to the previous inputs and previous outputs), such filters are more often described in the frequency-domain, by means of the \mathcal{Z} -transform [2]. For a given signal x , its \mathcal{Z} -transform is defined as:

$$X(z) = \sum_{k=-\infty}^{+\infty} x(k)z^{-k}. \quad (2)$$

This transform is an equivalent of the Laplace transform but for discrete-time signals and systems.

Applied to (1), we obtain a relationship between the \mathcal{Z} -transform of the input and the output, denoted U and Y respectively:

$$Y(z) = \sum_{i=0}^n b_i z^{-i} U(z) - \sum_{i=1}^n a_i z^{-i} Y(z) \quad (3)$$

and then $Y(z) = H(z)U(z)$, with

$$H(z) \triangleq \frac{\sum_{i=0}^n b_i z^{-i}}{1 + \sum_{i=1}^n a_i z^{-i}}. \quad (4)$$

H is called the *transfer function* of the filter, i.e. its input-output relationship in the z -domain. Its restriction to the unit circle ($\{z = e^{i\omega} \mid \forall \omega \in [0, 2\pi]\}$) is the Discrete-Time Fourier Transform (DTFT), so $H(e^{i\omega})$ gives the frequency response of the filter (if a sinusoid signal with pulsation ω_0 is applied as input of the filter, then the output will be a sinusoid with the same pulsation ω_0 , but amplified by $|H(e^{i\omega_0})|$, and dephased by $\arg(H(e^{i\omega_0}))$).

B. Filter specifications

In classical signal processing flow, filters are designed from specifications in frequency domain, in order to amplify (or preserve) signals in some frequency bands, and attenuate them in other bands.

A filter specification is then composed of several *passbands* (i.e. the gain of the filter for these frequencies should be bounded, often around 1) and *stopbands* (the gain should be lower than a given bound), formally described as:

$$\underline{\beta} \leq |H(e^{i\omega})| \leq \bar{\beta}, \quad \forall \omega \in [\omega_1, \omega_2]. \quad (5)$$

The lower bound $\underline{\beta}$ is equal to 0 for stopbands, and usually $\underline{\beta} = 1 - \delta$ and $\bar{\beta} = 1 + \delta$ for passbands.

For instance, Figure 1 exhibits a lowpass filter specification, mathematically described by

$$\begin{cases} 1 - \delta_p \leq |H(e^{i\omega})| \leq 1 + \delta_p, & \forall \omega \in [0, \omega_p] \quad (\text{passband}) \\ |H(e^{i\omega})| \leq \delta_s, & \forall \omega \in [\omega_s, \pi] \quad (\text{stopband}) \end{cases}$$




Fig. 1: A lowpass filter specification

For sake of generality, a filter specification will be described in this paper as a set of inequalities as in (5).

Sometimes, filter designers prefer to give the bounds in decibels (x dB means $10^{\frac{x}{20}}$ as a bound) and frequencies instead of pulsations (pulsation ω and frequency f are linked by $\omega = 2\pi \frac{f}{F_s}$ where F_s is the sampling frequency of the filter).

Due to Nyquist-Shannon theorem [2], it is only necessary to specify the frequency up to $\frac{F_s}{2}$ (or the pulsation up to π instead of 2π).

III. VERIFYING BOUNDS ON A TRANSFER FUNCTION

The purpose of this Section is to detail our method that verifies that the modulus of a transfer function H stays between two bounds $\underline{\beta}$ and $\bar{\beta}$ for all z taken on a segment of the unit circle, corresponding to a certain frequency band, i.e. $z = e^{j\omega}$ for all $\omega \in \Omega \subseteq [0, 2\pi]$. In the case when the given bounds cannot be verified, our intention is to compute approximations to problematic frequencies for which the bounds are violated.

We proceed in three steps. In Section III-A, we show how we can reduce the given problem to showing that a rational function with real coefficients stays between two bounds for real arguments taken in a subset of $[0, 1]$. In Section III-B, we then further reduce the problem to showing that a polynomial stays non-negative over a subset of $[0, 1]$. In Section III-C, we briefly describe our approach to computing problematic frequencies in the case when the verification does not succeed.

A. Reducing the problem to a real rational function

We wish to verify that $\underline{\beta} \leq |H(z)| \leq \bar{\beta}$ for all $z = e^{j\omega}$ with $\omega \in \Omega \subseteq [0, 2\pi]$. We suppose that H is given as a rational function $H(z) = \frac{b(z)}{a(z)}$ with real coefficients. Since we can suppose without lack of generality that $\underline{\beta} \geq 0$, this is equivalent to showing that

$$\underline{\beta}^2 \leq |H(z)|^2 \leq \bar{\beta}^2, \quad \forall z = e^{j\omega}, \omega \in \Omega. \quad (6)$$

Since $z = e^{j\omega}$ and the numerator and denominator polynomials a and b have real coefficients, conjugation of z yields $z^* = 1/z$ and conjugation of the polynomials has no effect. So we have

$$\begin{aligned} |H(z)|^2 &= \frac{b(z) b^*(z^*)}{a(z) a^*(z^*)} \\ &= \frac{b(z) b(1/z)}{a(z) a(1/z)} \\ &= \frac{v(z)}{w(z)}, \end{aligned} \quad (7)$$

where v and w also are polynomials with real coefficients, obtained by simplifying the fraction $\frac{b(z) b(1/z)}{a(z) a(1/z)}$.

We have hence reduced the problem to verifying that

$$\underline{\beta}^2 \leq \frac{v(z)}{w(z)} \leq \bar{\beta}^2, \quad \forall z = e^{j\omega}, \omega \in \Omega. \quad (8)$$

Taking now $t = \tan \frac{\omega}{2}$, we can write $z = e^{j\omega}$ as

$$z = e^{j\omega} = \cos \omega + j \sin \omega = \frac{1-t^2}{1+t^2} + j \frac{2t}{1+t^2}. \quad (9)$$

By formally composing v and w with the expression $z = \frac{1-t^2}{1+t^2} + j \frac{2t}{1+t^2}$, for example by formal evaluation with Horner's scheme, and clearing numerators and denominators, we can hence obtain four polynomials r , s , r_{im} , s_{im} , all with real coefficients, such that

$$|H(z)|^2 = \frac{v(z)}{w(z)} = \frac{v\left(\frac{1-t^2}{1+t^2} + j \frac{2t}{1+t^2}\right)}{w\left(\frac{1-t^2}{1+t^2} + j \frac{2t}{1+t^2}\right)} = \frac{r(t) + j r_{im}(t)}{s(t) + j s_{im}(t)}. \quad (10)$$

We can now observe that $|H(z)|^2$ is a real number and that the ratio $\frac{r(t)}{s(t)}$ is hence equal to the complex ratio $\frac{r(t) + j r_{im}(t)}{s(t) + j s_{im}(t)}$. We may therefore drop r_{im} and s_{im} . We have now reduced the problem to verifying that

$$\underline{\beta}^2 \leq \frac{r(t)}{s(t)} \leq \bar{\beta}^2, \quad \forall t = \tan \frac{\omega}{2}, \omega \in \Omega \subseteq [0, 2\pi], \quad (11)$$

where the both polynomials r and s have real coefficients and all other quantities, $\underline{\beta}^2, \bar{\beta}^2, \omega$ and t are all real numbers. We must hence no longer deal with complex ratios and complex numbers and have reduced the problem to verifying the bounds of a real rational function over an interval, subset of the reals.

Unfortunately, the mapping $t = \tan \frac{\omega}{2}$ maps the possible frequencies $\omega \in \Omega \subseteq [0, 2\pi]$ onto to the whole real axis. In our experiments, we found this difficult to handle, partly because the tool we used, Sollya, has very little support for unbounded intervals and partly because having unbounded intervals meant searching for the zeros of certain functions over such unbounded intervals, which we found numerically unstable (see Section III-C for more details). We hence apply a second mapping: $t = \frac{1-2\xi}{\xi(1-\xi)}$. Still by formally composing the polynomials r and s with the expression $t = \frac{1-2\xi}{\xi(1-\xi)}$, we obtain two polynomials p and q with real coefficients such that

$$|H(z)|^2 = \frac{r(t)}{s(t)} = \frac{p(\xi)}{q(\xi)}. \quad (12)$$

In the same step, we reduce the resulting rational function to its least terms to obtain $\frac{p(\xi)}{q(\xi)}$. In order to do so, we extended the tool we used, Sollya, with an algorithm to compute the gcd of two polynomials [3].

As the inverse mapping $\xi = \frac{t+2-\sqrt{t^2+4}}{2t}$ maps the reals onto the interval $[0, 1]$, we have hence reduced our problem to verifying that

$$\underline{\beta}^2 \leq \frac{p(\xi)}{q(\xi)} \leq \bar{\beta}^2, \quad \forall \xi \in \Xi \subseteq [0, 1]. \quad (13)$$

B. Verifying the bounds of a rational function by showing the non-negativity of a polynomial

In order to verify an instance of (13), we can suppose that the interval Ξ the arguments ξ vary in is not reduced to a point and that $\underline{\beta}^2 \neq \bar{\beta}^2$ (otherwise a simple evaluation or a check whether p and q are constant polynomials suffice). We can hence reduce the problem further to obtain:

$$-1 \leq \frac{2p(\xi) - (\bar{\beta}^2 + \underline{\beta}^2) q(\xi)}{(\bar{\beta}^2 - \underline{\beta}^2) q(\xi)} \leq 1, \quad \forall \xi \in \Xi. \quad (14)$$

Let

$$g(\xi) = 2p(\xi) - (\bar{\beta}^2 + \underline{\beta}^2) q(\xi)$$

and

$$h(\xi) = (\bar{\beta}^2 - \underline{\beta}^2) q(\xi).$$

It hence suffices to verify that

$$\frac{g(\xi)^2}{h(\xi)^2} \leq 1, \quad \forall \xi \in \Xi \quad (15)$$

which is equivalent to showing that

$$h(\xi)^2 - g(\xi)^2 \geq 0, \quad \forall \xi \in \Xi. \quad (16)$$

Let $f(\xi) = h(\xi)^2 - g(\xi)^2$. Again f is a polynomial with real coefficients. We have reduced our problem to showing that the value of this polynomial $f(\xi)$ stays non-negative over all $\xi \in \Xi \subseteq [0, 1]$, where the interval Ξ is easily obtained from the original frequency domain $\Omega = [\omega_1, \omega_2]$.

Our approach to showing that f stays non-negative over Ξ is similar to the one set out in [4]. We typically perform the following checks:

- (i) Check whether f is positive at some (arbitrarily chosen) point $\xi_1 \in \Xi$ by (interval arithmetic) evaluation of f at ξ_1 and that f has no zero over the whole interval Ξ . If so, f is non-negative over the whole interval Ξ .
- (ii) Check whether f is positive at both endpoints of the interval Ξ by (interval) evaluation at these endpoints and that it has exactly one zero over whole interval Ξ , not counting multiplicities. The zero it has in the interval hence is of even multiplicity and the polynomial stays non-negative over the whole interval.
- (iii) Check whether the interval Ξ can be split into subintervals such that one of the two aforementioned checks become satisfied.

We test whether a polynomial (with real coefficients) has no, one or more zeros over an interval, bounded subset of the reals, utilizing Sturm's theorems on the Sturm sequence of the polynomial, similarly as done in [4]. Sturm's theorem yields the number of real zeros of a real polynomial over a bounded interval, not counting multiplicities [5]. The tool we used, Sollya, includes a fast but rigorous implementation of Sturm's technique [6].

C. Numerically computing problematic frequencies

In the case when our checks verifying if a given transfer function $H(z)$ stays bounded in modulus by the two bounds $\underline{\beta}$ and $\bar{\beta}$ does not succeed, we numerically compute a list of problematic frequencies $\tilde{\omega}_i$, at which one of the bounds is violated. In contrast to the verification step which is completely rigorous in the sense that will never return a positive answer (i.e. the transfer function satisfies the given bounds) while the function actually does not, this numerical step is not fully rigorous. It may miss certain frequencies at which the bounds are violated. It is nevertheless pretty efficient with respect to speeding up the complete LTI filter verification algorithm we set out in Section IV-C, in particular concerning determining a reasonable verification margin (see below).

In our approach, we couple the verification process, described in Sections III-A and III-B, with the possibly needed step of computing problematic frequencies. These frequencies actually correspond, thru the different mappings $\omega \mapsto t \mapsto \xi$, to points ξ at which the polynomial eventually obtained, f , takes negative values. We determine these points as the (negative) extremum points of f . We therefore differentiate f and compute approximations to the zeros of f' by root isolation (still using Sturm's technique) and refinement with Newton-Raphson iterations. The tool we used, Sollya, offers all necessary basic bricks for these computations [6]. Once we obtain a list of points ξ_i at which f becomes negative, we remap these values ξ_i to a list of problematic frequencies, by following the inverse mappings $\xi \mapsto t \mapsto \omega$.




Fig. 2: If needed our algorithm returns the problematic frequencies as small intervals $\tilde{\omega}_i$

IV. COMPLETE ALGORITHM FOR A LTI REALIZATION VERIFICATION

The above-described algorithm can be used to verify whether the designed transfer function satisfies the band specifications, in other words assesses the quality of a design method. However, it is the quantization of filter's coefficients that

can drastically impact the behavior of the *implemented* filter, leading to violation of the initial band specifications. In this section we propose an approach on verifying an implemented filter against user-given band specifications.

To apply our algorithm from Section III, we need to first compute the transfer function that corresponds to the quantized coefficients. Just like polynomials, a rational transfer function can be evaluated in many different ways. Numerous filter structures have been developed over time. With some exceptions (Direct Forms) filter structures do not directly use the coefficients of transfer function but some modification of them. Deducing a transfer function from a structure differs from one structure to another [7], [8]. Moreover, they do not usually provide any error bound on the approximation.

In order to unify the approach of transfer function determination for any filter structure, we propose to use the Specialized Implicit Form (SIF) [9]. In this paper we restrict ourselves to the Single-Input Single-Output (SISO) case, but an extension to the Multiple-Input Multiple-Output (MIMO) case is straightforward. In Section IV-A we present a quick overview of this representation. In Section IV-B we propose our approach on the computation of the filter's transfer function with the corresponding bound on the approximation error. Finally, in Section IV-C we propose an algorithm for the verification of any implemented filter against some band specifications.

A. Specialized Implicit Form

The Specialized Implicit Form (SIF) has been proposed as an extension of the state-space form:

$$\begin{pmatrix} \mathbf{J} & \mathbf{0} & \mathbf{0} \\ -\mathbf{K} & \mathbf{I}_{n_x} & \mathbf{0} \\ -\mathbf{l} & \mathbf{0} & 1 \end{pmatrix} \begin{pmatrix} \mathbf{t}(k+1) \\ \mathbf{x}(k+1) \\ y(k) \end{pmatrix} = \begin{pmatrix} \mathbf{0} & \mathbf{M} & \mathbf{n} \\ \mathbf{0} & \mathbf{P} & \mathbf{q} \\ \mathbf{0} & \mathbf{r} & s \end{pmatrix} \begin{pmatrix} \mathbf{t}(k) \\ \mathbf{x}(k) \\ u(k) \end{pmatrix} \quad (17)$$

where $u(k)$ represents the input, and $y(k)$ the output, $\mathbf{x}(k+1)$ is the n_x stored states. Vector $\mathbf{t}(k+1)$ holds the n_t intermediate variables in the calculations on step k . These calculations are topologically ordered in the matrix \mathbf{J} , which is lower-triangular with 1s on the diagonal [10].

The diagonal matrix on the left side of the implicit equation (17) allows us to describe the sequence of computations within a filter. For example, $y \leftarrow \mathbf{m}_2(\mathbf{M}_1\mathbf{x})$ is computed as $\mathbf{t} \leftarrow \mathbf{M}_1\mathbf{x}$ and then $y \leftarrow \mathbf{m}_2\mathbf{t}$. So this sequence is described as

$$\begin{pmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{m}_2 & 1 \end{pmatrix} \begin{pmatrix} \mathbf{t} \\ y \end{pmatrix} = \begin{pmatrix} \mathbf{M}_1 \\ 0 \end{pmatrix} \mathbf{x}. \quad (18)$$

A similar approach as for (18) can be used to transform any linear signal processing algorithm or any input/output data-flow graph with delays, multiplications by constants and additions into a SIF. A systematic algorithm for the conversion of any linear data-flow graph represented in Simulink format to the SIF is presented in [10].

For example, a 3rd order linear digital filter implemented with a ρ -Direct Form II transposed (ρ DFII_t) algorithm [11], [12], will be described in Simulink as the data-flow graph




Fig. 3: Simulink data-flow of the ρ DFIlt form.

given in Figure 3. We can observe, that this structure uses only ten non-trivial coefficients. The corresponding SIF is the following 7×7 sparse matrix:

$$\mathbf{Z} \triangleq \begin{pmatrix} -\mathbf{J} & \mathbf{M} & \mathbf{n} \\ \mathbf{K} & \mathbf{P} & \mathbf{q} \\ \mathbf{l} & \mathbf{r} & \mathbf{s} \end{pmatrix} = \begin{pmatrix} \square & & & & & & \\ & \square & & & & & \\ & & \square & & & & \\ & & & \square & & & \\ & & & & \square & & \\ & & & & & \square & \\ & & & & & & \square \end{pmatrix}$$

where blue and pink rectangles are '-1' and '1', and circles are non-trivial coefficients.

Once rigorous filter analysis and implementation techniques are developed for the SIF formalism, they can be applied upon any linear realization (classical structures like direct forms, ρ -forms, lattice, etc. have already been directly converted to SIF [9], [13], [14]; others can be obtained from their input-output data flow graph [10]).

B. Reliable computation of the Transfer Function of a filter

The direct way to obtain the transfer function of a given SIF is to first convert it to a discrete-time state-space (dSS) representation $\mathcal{S} \triangleq (\mathbf{A}, \mathbf{b}, \mathbf{c}, d)$

$$\mathcal{S} \begin{cases} \mathbf{x}(k+1) = \mathbf{A}\mathbf{x}(k) + \mathbf{b}u(k) \\ y(k) = \mathbf{c}\mathbf{x}(k) + d u(k) \end{cases} \quad (19)$$

$$\mathbf{A} = \mathbf{K}\mathbf{J}^{-1}\mathbf{M} + \mathbf{P}, \quad \mathbf{b} = \mathbf{K}\mathbf{J}^{-1}\mathbf{n} + \mathbf{q} \quad (20a)$$

$$\mathbf{c} = \mathbf{l}\mathbf{J}^{-1}\mathbf{M} + \mathbf{R}, \quad d = \mathbf{l}\mathbf{J}^{-1}\mathbf{n} + s. \quad (20b)$$

The coefficients \mathbf{A} , \mathbf{b} , \mathbf{c} and d can be computed exactly from a SIF, since matrix \mathbf{J} is lower-triangular with 1s on the main diagonal.

Then, applying the \mathcal{Z} -transform of (2) to (19) we have the classical formula for the transfer function of a state-space [2]:

$$H(z) = \mathbf{c}(z\mathbf{I} - \mathbf{A})^{-1}\mathbf{b} + d. \quad (21)$$

By considering the eigendecomposition $\mathbf{V}\mathbf{E}\mathbf{V}^{-1}$ of the matrix \mathbf{A} (with $\{\lambda_i\}_{1 \leq i \leq n_x}$ its associated eigenvalues, i.e. the diagonal elements of \mathbf{E}), we have

$$H(z) = \mathbf{c}(z\mathbf{I} - \mathbf{V}\mathbf{E}\mathbf{V}^{-1})^{-1}\mathbf{b} + d \quad (22)$$

$$= \mathbf{c}\mathbf{V} \begin{pmatrix} \frac{1}{z-\lambda_1} & & & \\ & \ddots & & \\ & & \frac{1}{z-\lambda_n} & \\ & & & \ddots \end{pmatrix} \mathbf{V}^{-1}\mathbf{b} + d, \quad (23)$$

Therefore, the transfer function can be written as a rational function $H(z) = \frac{b(z)}{a(z)}$ with

$$a(z) = \prod_{j=1}^{n_x} (z - \lambda_j) \quad (24)$$

$$b(z) = \sum_{i=1}^{n_x} (\mathbf{c}\mathbf{V})_i (\mathbf{V}^{-1}\mathbf{b})_i \prod_{j \neq i} (z - \lambda_j) + d a(z) \quad (25)$$

Then, the accuracy of the computation of the transfer function relies on the accuracy of the computation of its eigendecomposition and on the computation of the polynomial coefficients. Exhibiting a bound on the accuracy of the eigendecomposition is a non-trivial task. Moreover, as it will be shown below, we may need to iteratively increase the accuracy of the computed transfer function. Thus, we propose to use Multiple Precision Floating-Point arithmetic to compute an approximation $\hat{H}(z)$ using formulas (24) and (25).

Suppose we have available a multiple precision eigensolver such that its error monotonically decreases and multiple precision basic bricks for matrix arithmetic. Then, increasing the precision of the computation of \hat{H} decreases its error $\|H - \hat{H}\|$ for any transfer function norm $\|\cdot\|$.

In order to bound this absolute error we propose the following:

Step 1: since we cannot exactly compute H , we first compute the approximation \hat{H} on the transfer function of the dSS $\mathcal{S} = (\mathbf{A}, \mathbf{b}, \mathbf{c}, d)$ using an eigendecomposition and Multiple Precision arithmetic (equations (24) and (25));

Step 2: we compute the system $\hat{\mathcal{S}}$ which exactly corresponds to the approximation $\hat{H}(z)$ using the controllable canonical state-space [2]: if $\{\hat{b}_i\}_{0 \leq i \leq n}$ and $\{\hat{a}_i\}_{1 \leq i \leq n}$ are the coefficients of the polynomial defining \hat{H} , then $\hat{\mathcal{S}}$ is given by $(\hat{\mathbf{A}}, \hat{\mathbf{b}}, \hat{\mathbf{c}}, \hat{d})$ with:

$$\hat{\mathbf{A}} = \begin{pmatrix} -\hat{a}_1 & 1 & & \\ \vdots & & \ddots & \\ \vdots & & & 1 \\ -\hat{a}_n & 0 & \dots & 0 \end{pmatrix}, \quad \hat{\mathbf{b}} = \begin{pmatrix} \hat{b}_1 - \hat{a}_1\hat{b}_0 \\ \vdots \\ \hat{b}_n - \hat{a}_n\hat{b}_0 \end{pmatrix} \quad (26)$$

$$\hat{\mathbf{c}} = (1 \ 0 \ \dots \ 0), \quad \hat{d} = \hat{b}_0$$

Step 3: we compute exactly the state-space difference $\Delta\mathcal{S} = \mathcal{S} - \hat{\mathcal{S}}$, which is the difference between outputs of the two state-spaces, using the classical formula [1]:

$$\Delta\mathbf{A} = \begin{pmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{A}} \end{pmatrix}, \quad \Delta\mathbf{b} = \begin{pmatrix} \mathbf{b} \\ \hat{\mathbf{b}} \end{pmatrix} \quad (27a)$$

$$\Delta\mathbf{c} = (\mathbf{c} \ -\hat{\mathbf{c}}), \quad \Delta d = d - \hat{d} \quad (27b)$$

Step 4: then, we can compute a bound on the approximation error $H - \hat{H}$. By definition of the state-space $\Delta\mathcal{S}$, we have $H = \hat{H} + \Delta H$ where ΔH is the transfer function of $\Delta\mathcal{S}$. Figure 4 illustrates the relationships between H , \hat{H} , ΔH and their corresponding state-space systems.

Using the relations between the transfer function and impulse response norms of a filter [15], we have

$$\forall \omega \in [0, 2\pi], \quad |\Delta H(e^{j\omega})| \leq \sup_{\omega} |\Delta H(e^{j\omega})| \leq \|\Delta h\|_1$$




Fig. 4: Dotted line: inexact transformation. Straight line: exact transformation.

where $\Delta h(k)$ is the impulse response of the system (response to the filter ΔS with transfer function ΔH to an impulse signal), and $\|\Delta h\|_1$ its ℓ_1 -norm, i.e. $\sum_{k=0}^{\infty} |\Delta h(k)|$. As shown in [16], [17], the norm $\|\Delta h\|_1$ can also be computed using the Worst-Case Peak-Gain (WCPG) measure of ΔS , denoted $\langle\langle \Delta S \rangle\rangle$, and defined by:

$$\langle\langle \Delta S \rangle\rangle \triangleq \sum_{k=0}^{\infty} |\Delta c(\Delta A)^k \Delta b| + |\Delta d|. \quad (28)$$

Using the result given in [17], we can compute $\langle\langle \Delta dSS \rangle\rangle$ with arbitrary precision. Finally, the approximation error $H(z) - \hat{H}(z)$ is bounded, for any z in the unit circle:

$$\forall \omega \in [0, 2\pi], \quad |H(e^{j\omega}) - \hat{H}(e^{j\omega})| \leq \Theta \quad (29)$$

where $\Theta = \langle\langle \Delta S \rangle\rangle + \varepsilon$ is the Worst-Case Peak Gain of the system ΔS computed with arbitrary small absolute error bounded by the $\varepsilon > 0$.

This way of first computing a Floating-Point approximation and then doing an exact inverse transformation to rigorously bound the error was inspired by S. Rumps verified inclusion techniques [18].

C. Reliable verification of the implemented filter

Now, given an implemented filter, we represent it in SIF and then compute, with some precision, an approximation \hat{H} on its transfer function along with a rigorous upper bound Θ on the corresponding $\langle\langle \Delta S \rangle\rangle$. We choose the precision of the computations based on an heuristic which increases the compute precision in a loop and is, hence, reasonably fast and accurate but does not provide any guarantee that the precision will eventually be enough.

Then, to verify whether the exact magnitude response evaluated on the unit circle is in the interval $[\underline{\beta}; \bar{\beta}]$, it is sufficient that the approximation $|\hat{H}(e^{j\omega})|$ for $\omega \in \Omega$ is in the interval $[\underline{\beta} + \Theta; \bar{\beta} - \Theta]$. See Figure 5 for an illustration.




Fig. 5: Verify whether $|\hat{H}(e^{j\omega})|$ is in $[\underline{\beta} + \Theta; \bar{\beta} - \Theta]$

If the verification is not successful, we obtain a list of problematic frequencies with algorithm from Section III. Thus, we can compute the maximum excess of the bounds, enlarge the band by this amount and repeat the process. For example, on Figure 6i we suppose that the approximation $|\hat{H}(e^{j\omega})|$ is too close to the bound or Θ is too large. Then, we enlarge the band margin up to some new bound $\bar{\beta}'$ while decreasing Θ , and perform the verification again with the updated band.

This algorithm may yield a false negative answer if the initial precision of the transfer function computation was not large enough, which we never observed on the numerical examples we ran the algorithm on.

The detailed description of the heuristics that we used can be found in the long version of the paper that will be published along with our software implementation.




Fig. 6: If verification fails (i), we enlarge the bound (ii)

V. NUMERICAL EXAMPLES

The algorithm presented in Section III was implemented using a modified Sollya tool³ and further bound with our automatic filter generator using pythonSollya⁴. We use an implementation of the algorithm for the WCPG⁵ computation in arbitrary precision [17], which is written in C, using GNU MPFR version 3.1.12, GNU MPFI version 1.5.1 and CLAPACK version 3.2.1. Experiments were done on a laptop computer with an Intel Core i5 processor running at 2.8 GHz and 16 GB of RAM.

In this section we consider four simple band specifications: lowpass, highpass, bandstop and bandpass filters, with sampling frequency $F_s = 48\text{kHz}$. Band specifications are presented schematically on the Figure 7.

The two major state-of-the-art tools for the filter design are Matlab Filter Design Toolbox and Python SciPy. We consider three widely-used IIR filter design methods: Butterworth [19], Elliptic [20], [21] and Chebyshev [22].

Our approach to verifying a transfer function against band specification can be used in several user-cases.

The first one is an *a posteriori* verification of rational approximation on the transfer function. In this setting, our algorithm permits to assess the quality of design method. Then our goal is now to verify, using the algorithm from Section IV-C, the transfer functions obtained with the aforementioned methods

³Git branch gcd at <https://scm.gforge.inria.fr/anonscm/git/sollya/sollya.git>

⁴<https://scm.gforge.inria.fr/anonscm/git/metalibm/pythonSollya.git>

⁵<https://scm.gforge.inria.fr/anonscm/git/metalibm/wcpG.git>

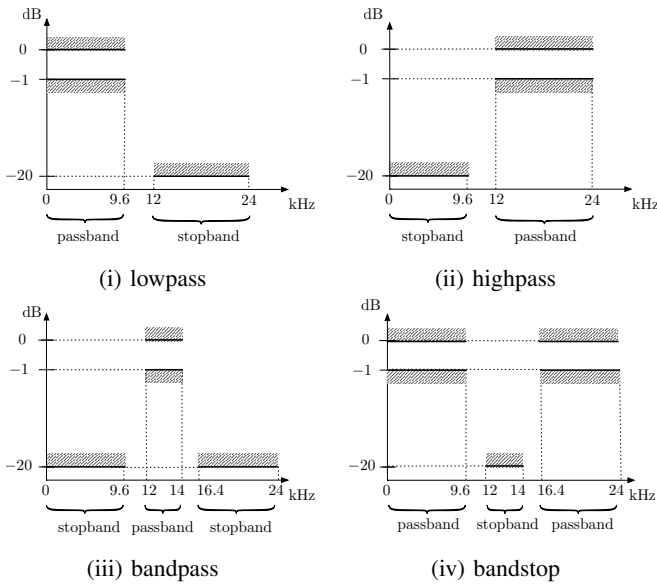


Fig. 7: Band specifications for numerical examples

and tools. If the verification is not successful, a safe margin for the band specifications is provided. As we can observe in Table I, only some transfer functions verify the band specifications. On the other hand, since the design tools are implemented in double precision Floating-Point arithmetic, a margin of order 10^{-16} may seem acceptable. However, we can observe that a relatively large margin 0.0448dB is required (according to our algorithm) for the majority of transfer functions designed with the Elliptic method. This information can permit the filter designer to decide which method to use, and whether the computed safe margin is acceptable or not. Overall, we can conclude that Matlab proposes relatively better transfer function designs, in comparison to the Python SciPy. Thus, our algorithm can be used to verify the quality of a filter design method.

		Butterworth	Chebyshev	Elliptic
		margin (dB)	margin (dB)	margin (dB)
lowpass	Matlab	1.29e-17	7.93e-17	✓
	SciPy	2.14e-15	4.48e-2	4.48e-2
highpass	Matlab	2.77e-16	6.94e-17	4.48e-2
	SciPy	3.02e-15	2.29e-16	4.48e-2
bandpass	Matlab	3.04e-17	✓	✓
	SciPy	✓	4.48e-2	4.48e-2
bandstop	Matlab	4.59e-16	3.09e-15	✓
	SciPy	✓	6.36e-15	7.02e-6

TABLE I: Comparison of Matlab and SciPy filter designs

More importantly, the second use-case for our algorithm is an *a posteriori* verification of an *implemented* filter. Under implemented filter we consider a transfer function realized with some structure and with its coefficients quantized to some Fixed-Point Formats. Different filter structures have different sensitivity with respect to quantization of its coefficients, so their finite precision behavior may be different. On top

of that, usually the structure coefficients are computed out of the transfer function using double-precision floating point arithmetic, which introduces additional rounding errors into the actually implemented coefficients. For example, the state-space controllable form is based directly on the transfer function coefficients, while the balanced state-space system is obtained by performing numerous similarity transformations of the controllable form.

We design two transfer functions via Matlab: by Butterworth and Chebyshev methods. We discard the Elliptic method since it requires a large design margin for a designed transfer function even before its Fixed-Point implementation.

The target Fixed-Point implementations have coefficients quantized to 32, 16 and 8 bits.

Consider the following filter structures:

- Direct Form II transposed (DFII_t): a classic structure, with large output error
- ρ Direct Form II transposed (ρ DFII_t): uses a weighted delay operator, has $2n$ additional parameters enabling structure optimization, small output error
- Controllable State-Space: obtained directly from the TF coefficients, reasonable output error
- State-Space: a classic structure for control systems, reasonable output error
- Balanced State-Space [23]: an optimized State-Space structure that minimizes the coefficient sensitivity to quantization, small output error

We illustrate the comparison between different filter implementation schemes for the simple lowpass filter in Figure 7. We can observe that the safe margin on the band specification, as expected, increases with the quantization of the coefficients. Exceptions similar to the Balanced State-Space may occur: in double-precision, its transfer function does not pass the verification but the coefficient quantization to 16 and 8 bits “luckily” moves the transfer function of the systems in the right direction. On the other hand, we can observe that even though the DFII_t yields very large output error in the time-domain in comparison to the ρ DFII_t [11], [12], it performs better with respect to the frequency domain specification. This is due to the fact that DFII_t uses directly the transfer function’s coefficients, while ρ DFII_t performs various optimizations of the weight ρ to reduce the output error. Finally, for the cases when the transfer function respects the band specifications our algorithm quickly gives an answer. It computes the safe margin, especially when an overflow is on the edge of the band that may take significant amount of time.

VI. CONCLUSION AND PERSPECTIVES

With this paper, a rigorous method to verify a transfer function of LTI digital filters against band specifications in the frequency domain has been developed. It relies on translating the problem of verifying bounds on magnitude response evaluated on a unit circle to the verification of positivity of a real polynomial. Our algorithm guarantees that no false positives occur. In the case of unsuccessful verification a list of problematic frequencies is provided, for which we compute the

	filter	Butterworth			Chebyshev		
		wordlength	32	16	8	32	16
DFIIt	margin(dB)	✓	✓	1.17e-2	1.97e-10	7.03e-6	7.42e-3
	time	12s	15s	1min17s	1min22s	35s	18s
ρ DFIIt	margin(dB)	✓	6.78e-2	✗	1.21e-9	4.72e-4	8.35e-2
	time	13s	3min17s	unstable	1min26s	1min02s	17s
State-Space	margin(dB)	✓	9.11e-7	9.64e-3	2.06e-10	7.64e-6	2.23e-3
	time	15s	2min4s	1min01s	23s	49s	18s
State-Space	margin(dB)	6.77e-11	✓	✓	3.71e-11	3.38e-6	4.25e-4
Balanced	time	6min27s	11s	10s	1min15s	1min	1min08s
State-Space	margin(dB)	✓	9.11e-7	9.64e-3	2.06e-10	7.63e-6	2.23e-3
Controllable	time	15s	2min05s	1min02s	24s	50s	19s

TABLE II: Comparison of implemented structures based on a MATLAB generated lowpass filter.

maximum overflow over the band specification. We propose an implementation using a combination of interval and rational arithmetic in Sollya tool.

We applied this method to develop an approach on the verification of any implemented filter, which is possible thanks to our automatic filter code generator. It relies on the multiple precision computation of a transfer function corresponding to the filter with quantized coefficients. Then, we bound the approximation error using the Worst-Case Peak Gain measure in arbitrary precision.

This approach opens various possibilities on the verification of digital filters, as well as the transfer function design tools. Even naive comparing the state-of-the-art tools and methods revealed weaknesses in the Python SciPy library, as well as frequent issues with the Elliptic method for IIR filter design in both Scipy and Matlab.

Moreover, our verification algorithm can be applied to compare different filter structures with various Fixed-Point settings. Such a comparison offers a filter designer an overview of the implementation possibilities. On top of that, the information on the safe margin can be used to correct a design that fails to verify the initial band specifications. For instance, we can narrow down the initial band by the safe margin and re-design the filter with the new band specifications. Then, the re-computed filter with quantized coefficients should respect the initial conditions. Such an approach can be also applied to improve the existing rational approximation algorithms [24].

However, overall time-efficiency of our implementation can still be improved. Passing too much time on the computation of the safe margin can be a significant drawback were the algorithm to be used during the exploration of a large design space. Nonetheless, this limitation can be overcome by setting an initial acceptable design margin, which directly depends on the application of the filter.

Additionally, it would be interesting to study the connection and equilibrium of the errors in the frequency and time domains to obtain optimal and stable designs.

REFERENCES

- [1] T. Kailath, *Linear Systems*. Prentice-Hall, 1980.
- [2] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*, 3rd ed. NJ, USA: Prentice Hall Press, 2009.
- [3] B. W. Char, K. O. Geddes, and G. H. Gonnet, "Godhe: Heuristic polynomial gcd algorithm based on integer gcd computation," *Journal of Symbolic Computation*, vol. 7, no. 1, pp. 31 – 48, 1989.
- [4] S. Chevillard, J. Harrison, M. Joldes, and C. Lauter, "Efficient and accurate computation of upper bounds of approximation errors," *Theoretical Computer Science*, vol. 412, no. 16, pp. 1523 – 1543, 2011.
- [5] M.-F. Roy, *Basic algorithms in real algebraic geometry and their complexity: from Sturm's theorem to the existential theory of reals*, ser. Expositions in Mathematics. de Gruyter, 1996, vol. 23, in F. Broglia (Ed.), *Lectures in Real Geometry*.
- [6] S. Chevillard, C. Lauter, and M. Joldes, *Users' manual for the Sollya tool, Release 6.0*, LIP, LIP6, LORIA, CNRS.
- [7] A. Varga, "Computation of transfer-function matrices of generalized state-space models," *Int Journal of Control*, vol. 50, no. 6, 1989.
- [8] G. A. Constantinides, P. Y. K. Cheung, and W. Luk, "Wordlength optimization for linear digital signal processing," *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 22, no. 10, pp. 1432–1442, 2003.
- [9] T. Hilaire, P. Chevrel, and J. Whidborne, "A unifying framework for finite wordlength realizations," *IEEE Trans. on Circuits and Systems*, vol. 8, no. 54, pp. 1765–1774, 2007.
- [10] T. Hilaire, A. Volkova, and M. Ravoson, "Reliable fixed-point implementation of linear data-flows," in *Proc. IEEE Workshop on Signal Processing Systems (SiPS)*, 2016.
- [11] Z. Zhao and G. Li, "Roundoff noise analysis of two efficient digital filter structures," *IEEE Transactions on Signal Processing*, vol. 54, no. 2, pp. 790–795, February 2006.
- [12] G. Li and Z. Zhao, "On the generalized DFIIt structure and its state-space realization in digital filter implementation," *IEEE Trans. on Circuits and Systems*, vol. 51, no. 4, pp. 769–778, April 2004.
- [13] T. Hilaire and P. Chevrel, "Sensitivity-based pole and input-output errors of linear filters as indicators of the implementation deterioration in fixed-point context," *EURASIP Journal on Advances in Signal Processing*, January 2011.
- [14] A. Volkova and T. Hilaire, "Fixed-point implementation of lattice wave digital filter: Comparison and error analysis," in *23rd European Signal Processing Conference (EUSIPCO)*, Aug 2015, pp. 1118–1122.
- [15] S. P. Boyd and J. Doyle, "Comparison of peak and rms gains for discrete-time systems," *Syst. Control Lett.*, vol. 9, no. 1, pp. 1–6, June 1987.
- [16] V. Balakrishnan and S. Boyd, "On computing the worst-case peak gain of linear systems," *Systems & Control Letters*, vol. 19, 1992.
- [17] A. Volkova, T. Hilaire, and C. Lauter, "Reliable evaluation of the worst-case peak gain matrix in multiple precision," in *Computer Arithmetic (ARITH)*, 2015 IEEE 22nd Symposium on, June 2015, pp. 96–103.
- [18] S. M. Rump, "Reliability in computing: The role of interval methods in scientific computing," R. E. Moore, Ed. Academic Press, 1988, ch. Algorithms for Verified Inclusions Theory and Practice, pp. 109–126.
- [19] S. Butterworth, "On the theory of filter amplifiers," *Wireless Engineer*, vol. 7, pp. 536–541, 1930.
- [20] E. I. Zolotarev, "Application of elliptic functions to problems about functions with least and greatest deviation from zero," *Zap. Imp. Akad. Nauk. St. Petersburg*, vol. 30, no. 5, 1877, (in Russian).
- [21] W. Cauer, "Synthesis of linear communication networks," *McGraw-Hill, New York*, 1958.
- [22] R. W. Daniels, *Approximation Methods for Electronic Filter Design*. New York: McGraw-Hill., 1974.
- [23] M. Gevers and G. Li, *Parametrizations in Control, Estimation and Filtering Problems*. Springer-Verlag, 1993.
- [24] I. W. Selesnick, M. Lang, and C. S. Burrus, "Magnitude squared design of recursive filters with the chebyshev norm using a constrained rational remez algorithm," in *Proceedings of IEEE 6th Digital Signal Processing Workshop*, Oct 1994, pp. 23–26.