

## A certified infinite norm for the validation of numerical algorithms

Sylvain Chevillard,
Christoph Lauter


# A certified infinite norm for the validation of numerical algorithms 

Sylvain Chevillard, Christoph Lauter

Décembre 2006


#### Abstract

The development of numerical algorithms requires the bounding image domain of functions, in particular functions $\varepsilon(x)$ associated to an approximation error. This problem can often be reduced to computing the infinite norm $\|\varepsilon(x)\|_{\infty}$ of the given function $\varepsilon(x)$. For instance, the development of elementary function operators in hard- and software makes use of such algorithms. Implementations for computing in practice highly accurate floating-point approximations to infinite norms are known and available. Nevertheless, no highly precise, sufficiently fast and certified or self-validating algorithms are available. Their results could be seen as an element in the correctness proof of safety critical or provenly guaranteed implementations. We present an algorithm for computing infinite norms in interval arithmetic. The algorithm is optimized for functions representing absolute or relative approximation errors that are ill-conditioned because of high cancellation. It can handle even functions that are numerically unstable on floating-point points because they are defined there only by continuous extension. In addition the given algorithm is capable of generating a correctness proof for an infinite norm instance by retaining its computational tree.


Keywords: infinite norm, optimization, interval arithmetic, certified algorithm, error analysis, approximation error

## Résumé

Le développement d'algorithmes numériques nécessite de borner certaines fonctions, en particulier les fonctions représentant une erreur d'approximation. Ce problème se réduit au calcul de la norme infinie $\|\varepsilon(x)\|_{\infty}$ de la fonction d'erreur $\varepsilon(x)$. Par exemple, le développement de fonctions élémentaires, tant au niveau logiciel que matériel, utilise ce genre de calcul.
Il existe déjà des implémentations de la norme infinie fournissant une très bonne approximation de la valeur réelle de la norme. Cependant, il n'existe pas d'algorithme capable de fournir un résultat à la fois précis et sûr. On entend par sûr, un algorithme qui renvoie une valeur majorant la norme réelle et qui fournit par ailleurs un certicat prouvant la validité de cette majoration.
Nous proposons un algorithme de calcul de la norme infinie utilisant l'arithmétique d'intervalles. Cet algorithme est optimisé pour les fonctions correspondant à une erreur relative ou absolue, c'est-à-dire des fonctions numériquement très mal conditionnée du fait d'importantes cancellations. Notre algorithme peut aussi, dans une certaine mesure, travailler avec des fonctions numériquement instables à proximité de certains points où elles ne sont définies que par continuité. Enfin, notre algorithme peut retenir l'arbre des calculs qu'il a effectués afin de produire une preuve de correction du résultat de son calcul.

Mots-clés: norme infinie, optimisation, arithmétique d'intervalles, algorithme certifié, analyse d'erreur, erreur d'approximation

## 1 Introduction

The development of a numerical algorithm, such as scientific code[8, 6], elementary function implementation[1] or control applications, consists generally of three steps. Firstly, the given problem is expressed as a mathematical model. This mathematical model may still make usage of high level concepts or functions that are not directly supported by current combinations of processors, programming languages etc. such as combined, non-elementary functions. Secondly, the mathematical model is simplified to match already more closely the available hardware and software system. In this step, approximations take place. For example, transcendental functions may be approximated by rational functions. Or combined functions may be replaced by a combination of approximations of elementary functions. In a third step, the simplified and approximated rational mathematical model is implemented in a floating-point environment, provided, for example, by the IEEE 754 standard[2]. Floating-point numbers generally are finite dyadic approximations to real numbers in a finite range around zero.

All three steps of modelling a given problem imply errors. The mathematical model does not exactly match reality. Simplified to a rational model, is it subject to approximation errors. Finally, floating-point computations induce round-off error. In order to be certain of the significance of a numerical result, quantities appearing in the given model must be bounded. Computed values must be proven to be contained in the finite range of the floatingpoint environment. In addition, approximation and round-off errors must be shown to be less than a priori specified bound.

Round-off errors induced by a floating-point arithmetic are discrete, non-analytical, discontinuous functions of the inputs of the different basic arithmetical operators. Their bounding has been studied for example in [6]. Automatic tight boundings can be computed and proven using for example the Gappa* tool[4]. In this article we will not further consider them.

Values and approximation errors in an given model can be considered as almost everywhere continuous functions $\varepsilon: \mathbb{R}^{n} \mapsto \mathbb{R}$ of the inputs. Bounding them means computing their extrema in a given domain $I \subset R^{n}$. If quantities, especially errors, are mostly symmetric or strictly positive or negative, sufficient bounding may be acheived by computing the infinite (or infinity) norm (infnorm) of the function $\varepsilon$ defined as

$$
\|\varepsilon(\mathbf{x})\|_{\infty}^{I}=\sup _{\mathbf{x} \in I}|\varepsilon(\mathbf{x})|
$$

Computing the infinite norm of a function $\varepsilon$, given as a expression or a numerical operator, is for itself a numerical problem. High-quality, approximate, floating-point solutions to the infinite norm computational problem exist. General techniques and considerations are described in [10]. In the case where $\varepsilon$ is a multivariate function, computing a infinite norm is a particular case of global optimization. In this article we will consider only univariate functions. We attract the reader's attention to [5] concerning the multivariate case.

Tools like Maple ${ }^{\dagger}$ or Matlab ${ }^{\ddagger}$ implement general purpose numerical approximation algorithms for computing an infinite norm of an univariate function. Both algorithms are not clearly specified in terms of the quality of the returned approximation. Matlab uses hardware, i.e. IEEE 754 double, precision and is hence limited to well-conditioned infnorm problems. The infnorm algorithm in Maple's numapprox package tends to provide overestimations of

[^0]the infinite norm's true value but can be shown also to return underestimations on some particular functions.

Such approximate solutions may be sufficient in the development phase of a numerical implementation. But whenever it comes to prove its correctness, in particular, if the implementation is safety critical, numerical approximations do no longer suffice. An algorithm that provides a certified or self-validating result such as an interval with guarenteed lower and upper bounds of the computed approximation of the infinite norm is needed here.

The authors' work has been motivated by infinite norm problems in the development of correct rounding transcendental elementary functions such as $e^{x}, \log x, \sin x$. The correct rounding, i.e. bit-exact result, correctness proof mainly relies on showing a maximal approximation and round-off error bound[4]. Similar problems in the context of safety critical implementations of combined transcendental functions have been considered for example in[3]. The computation of well-specified approximations to infinite norms are also at the base of works like[11].

In this article we present an algorithm for computing a upper and a lower bound for infinite norms on univariate functions $\varepsilon(x) \in \mathbb{C}^{2}$ in self-validating and hence certified way. The algorithm is especially optimized for functions $\varepsilon(x)$ that are ill-conditioned because of cancellation and numerically unstable at some floating-point numbers because they are defined only by continous extension in these points. The implementation of the algorithm is still under development and is integrated to a software tool ${ }^{\S}$.

This article is organized as follows: in the next section 2, we give the specifications of our algorithm and explain these design choices. In section 3.1 we give the algorithm as well as a correctness proof sketch. This general algorithm makes use of some particular interval arithmetic evaluation techniques that we present in section 3.2. These techniques are used in particular for bracketing the zeros of a function. Section 3.3 clarifies this point. Our algorithm is capable of retaining its computational tree for generating a proof of the generated result. The main considerations on this point are given in section 4 . Some examples in section 5 lead the reader to our conclusions in section 6 .

## 2 Specifications of our infnorm algorithm

Let be $f: \mathbb{R} \mapsto \mathbb{R}$ a function to be shown to be correctly implemented, i.e. approximated within a specified error bound. Let be $p: \mathbb{R} \mapsto \mathbb{R}$ the approximation to $f$ used in the implementation. So the absolute respectively relative approximation error of $p$ with regard to $f$ is a function $\varepsilon: \mathbb{R} \mapsto \mathbb{R}$ defined as $\varepsilon(x)=p(x)-f(x)$ respectively $\varepsilon(x)=\frac{p(x)-f(x)}{f(x)}$. In the framework of elementary function development, $f$ is a transcendental function and $p$ a polynomial with floating-point coefficients $[1,3]$. If $p$ and $f$ are continuous and continously differentiable functions that are not identically zero on no sub-interval and if $\varepsilon$ is finite everywhere (which is the case in pratical implementations), $\varepsilon$ is almost everywhere continuous and continously differentiable.

The first requirement our algorithm shall fullfil is implied by the fact that we want a certified result:

Requirement 1. The algorithm implementing the infinite norm of a function must always give an upper-bound of the real value of the infnorm of the function.

[^1]It would be possible to always answer $+\infty$ but this would be perfectly useless in practice. In order to estimate the order of error made by a certified infinite norm algorithm, a lower bound for the actual value is needed. This can be obtained with this second requirement:

Requirement 2. The algorithm shall give a lower-bound of the real value of the infnorm of the function. Thus if the orders of magnitude of the upper and the lower bounds are the same, it can be concluded that the result of the algorithm is accurate enough for the problem in consideration.

Since $p$ approximates $f$, the order of magnitude of $\varepsilon(x)$ is much lower than the order of magnitude of $p$ or $f$. In other words, the functions $\varepsilon(x)=p(x)-f(x)$ respectively $\varepsilon(x)=$ $\frac{p(x)-f(x)}{f(x)}$ are ill-conditioned due to the cancellation in the substraction $p(x)-f(x)$.

Even if we want the algorithm to handle also other functions as functions defined as $\varepsilon(x)=p(x)-f(x)$ or $\varepsilon=\frac{p(x)-f(x)}{p(x)}$, this observation leads to a third requirement to our algorithm:

Requirement 3. The algorithm shall take in input functions defined by an explicite expression tree. Ill-conditionned functions defined in this way shall be overcome by the usage of high intermediate precision and recorrelation[5, 3] techniques for interval arithmetic.

Let us still make one observation on the functions $\varepsilon(x)=\frac{p(x)-f(x)}{f(x)}$ we are especially interested in. Suppose that in the given domain $I=[a ; b]$ the infinite norm $\|\varepsilon(x)\|_{\infty}^{I}$ is to be computed on, $f(x)$ has a zero $z$, i.e. $f(z)=0$. If $p(z)=0$ in the same point $z$ and $\lim _{x \rightarrow z} p^{\prime}(x)-f^{\prime}(x)=c_{1}$ and $\lim _{x \rightarrow z} f^{\prime}(x)=c_{2}$ exist, $\varepsilon(z)=c=\frac{c_{1}}{c_{2}} \in \mathbb{R}$ is nevertheless well-defined in $z$ by continuous extension. In consequence $\|\varepsilon(x)\|_{\infty}^{I} \neq \infty$ as the pole in $z$ of $\varepsilon(z)$ might suggest and as might be computed by pure interval arithmetic.

We formulate thus the following additional requirement:
Requirement 4. If the expression tree for $\varepsilon(x)$ has some pole in the given input domain that may be extended by continuity, the algorithm for computing the infinite norm of $\varepsilon$ on $I$ shall return an upper bound different from $+\infty$.

In order to ensure that the algorithm respects its specifications, it should be carefully proven. However, the implementation could contain bugs. Moreover, some users of the algorithm want to provide proofs for the results of intermediate computations in the development of an algorithm they had been using an infinite norm algorithm for. This yields to a last requirement:
Requirement 5. The algorithm shall give, in addition to the result, a formal proof which can be checked externally and which ensures that the interval result is really bounding the mathematical infnorm value.

Naturally, we want our algorithm to be as performant as possible using the least memory possible.

## 3 The algorithm

We are going to present now our infinite norm algorithm. Let us remember that the algorithm requires the function $\varepsilon(x)$, given as an expression tree, to be at least $\mathcal{C}^{2}$ and formally differentiable.

In the following, if $\varepsilon$ is a function and $I$ an interval, $\varepsilon[I]$ will denote the set

$$
\varepsilon[I]=\{y \in \mathbb{R}, \exists x \in I, y=\varepsilon(x)\}
$$

It is well-known that this set is an interval when $\varepsilon$ is continuous.

### 3.1 General scheme of the algorithm

Basically, the algorithm is simple:
infnorm: input a function $\varepsilon \in \mathcal{C}^{2}(I)$ and a compact interval $I=[a, b]$ :

1. formally differentiate the function $\varepsilon$;
2. search a list of intervals $I_{1}, \cdots, I_{p}$ such that every zero of $\varepsilon^{\prime}$ lies in at least one of the $I_{k}$. Note that some $I_{k}$ may not contain any zero of $\varepsilon^{\prime}$ and reciprocally that two zeros may belong to the same $I_{k}$;
3. add $I_{0}=[a, a]$ and $I_{p+1}=[b, b]$ to the list;
4. compute $J_{0}, \cdots, J_{p+1}$ such that for each $k, \varepsilon\left[I_{k}\right] \subseteq J_{k}$.
5. compute the inner- and outer- enclosure of $g[I]$ from the intervals $J_{0}, \cdots, J_{p+1}$. This means compute intervals $I E$ and $O E$ such that

$$
\forall y \in I E, \exists x \in I, y=\varepsilon(x)
$$

and such that

$$
\forall k, \forall y \in J_{k}, y \in O E
$$

The computation of this enclosures will be explained below.
6. Return

$$
\left[\max \left\{\left|I E_{\ell}\right|,\left|I E_{r}\right|\right\}, \max \left\{\left|O E_{\ell}\right|,\left|O E_{r}\right|\right\}\right]
$$

as an interval containing $\|\varepsilon\|_{\infty}^{I}$
Since the interval $I=[a, b]$ is compact and $\varepsilon$ is continuous, we know that $\varepsilon$ reaches its minimum at a point $x_{m}$ and its maximum at a point $x_{M}$. Since $\varepsilon$ is differentiable, $x_{m}$ is either a bound of the domain ( $a$ or $b$ ) or $\varepsilon^{\prime}\left(x_{m}\right)=0$. The same holds for $x_{M}$. It follows that both $x_{m}$ and $x_{M}$ are in some interval of the list $I_{0}, \cdots, I_{p+1}$.

Lemma 3.1. $\varepsilon\left(x_{m}\right)$ does not belong to the interior of IE. The same thing holds for $x_{M}$.
Proof. Suppose that $\varepsilon\left(x_{m}\right)$ belongs to the interior of $I E$. Thus, there exists some $z \in I E$ such that $z<g\left(x_{m}\right)$. By applying the definition of $I E$, we would have some $x$ such that $\varepsilon(x)=z$ yielding contradiction. Hence, since $\varepsilon\left(x_{m}\right)$ is minimal and $\varepsilon(x)=z<\varepsilon\left(x_{m}\right)$. The proof is the same for $x_{M}$.

Lemma 3.2. $\varepsilon\left(x_{m}\right) \in O E$. The same holds for $x_{M}$.
Proof. Since $x_{m}$ belongs to one of the $I_{k}, \varepsilon\left(x_{m}\right)$ belongs to $\varepsilon\left[I_{k}\right]$ and then $\varepsilon\left(x_{m}\right) \in J_{k}$. Applying the definition of $O E, \varepsilon\left(x_{m}\right) \in O E$. The proof is the same for $x_{M}$.

Figure 1: The outer enclosure of the $J_{k}$.


Let us see now how to compute $I E$ and $O E$. For $O E$ we take the convex enclosure of the union of the $J_{k}$ which is defined by $O E_{\ell}=\min \left\{J_{0 \ell}, \cdots, J_{(p+1) \ell}\right\}$ and $O E_{r}=\max \left\{J_{0 r}, \cdots, J_{(p+1) r}\right\}$ where $O E=\left[O E_{\ell}, O E_{r}\right]$. It trivially satisfies the required property for $O E$.

For $I E$ we take $] I E_{\ell}, I E_{r}\left[\right.$ where $I E_{\ell}=\min \left\{J_{0 r}, \cdots, J_{(p+1) r}\right\}$ and $I E_{r}=\max \left\{J_{0 \ell}, \cdots, J_{(p+1) \ell}\right\}$ where the subscript $\ell$ denotes the lower bound of an interval and the subscript $r$ denotes the upper bound.

Lemma 3.3. The previous way of computing IE actually provides an inner enclosure as defined above.

Proof. Let $I E_{\ell} \leq y \leq I E_{r}$. Then $y \geq \min \left\{J_{0 r}, \cdots, J_{(p+1) r}\right\}$. Let $k$ be the index for which the minimum is reached: $y \geq J_{k r}$. Since $f\left[I_{k}\right] \subseteq J_{k}$ there exists $u$ such that $f(u) \leq J_{k r} \leq y$. With the same argument, there exists $v$ such that $y \leq f(v)$. Since $f[I]$ is a interval, $y \in f[I]$ and then $\exists x \in I, y=f(x)$ which is the required property.

The computation of $I E$ and $O E$ from the $J_{k}$ can be performed incrementally as the $J_{k}$ are calculated.

It is clear that $I E \subseteq O E$ and, hence, $O E \backslash I E$ is made of the union of two intervals: $\left[O E_{\ell}, I E_{\ell}\right]$ and $\left[I E_{r}, O E_{r}\right]$. By the lemmata 3.1 and 3.2 the minimum and maximum of $g$ lie in these intervals. It follows that

$$
\|\varepsilon\|_{\infty}^{I} \in\left[\max \left\{\left|I E_{\ell}\right|,\left|I E_{r}\right|\right\}, \max \left\{\left|O E_{\ell}\right|,\left|O E_{r}\right|\right\}\right]
$$

This is the value returned by the algorithm given above.

Figure 2: The (right bound of the) inner enclosure of the $J_{k}$.


Let us now show how the $I_{k}$ are found and in which way the $J_{k}$ are computed out of them. Clearly if every $J_{k}$ were exactly equal to $\varepsilon\left[I_{k}\right], O E$ would precisely be equal to $\left[\varepsilon\left(x_{m}\right), \varepsilon\left(x_{M}\right)\right]$. Let $I_{k}$ denote the interval containing $x_{M}$. If $J_{k}$ is affected by arithmetical errors, we have $\varepsilon\left[I_{k}\right] \subset J_{k}$ and thus $O E_{r}$ is hence greater or equal to $J_{k r}$. It follows that the more precise the $J_{k}$ are, the more precise $O E$ will be. Here the term precise stands for some measure on an interval $A$ with regard to another interval $A^{\prime} \supseteq A$ it approximates. A measure for this phenomenon could be the sum of the relative diameters of the union of intervals in $A \backslash A^{\prime}$ with regard to the diameter of $A$.

Let $I_{k}$ be an interval such that $x_{M} \in I_{k}$. If $I_{k}$ is perfectly precise (that is if $I_{k}=\left[x_{M}, x_{M}\right]$ ), we have $\varepsilon\left[I_{k}\right]=\left[\varepsilon\left(x_{M}\right), \varepsilon\left(x_{M}\right)\right]$. But if $I_{k}$ is wider, $\varepsilon\left[I_{k}\right]$ will be of the form $\left[u, \varepsilon\left(x_{M}\right)\right]$ with $u$ the smaller as $I_{k}$ becomes the wider. Thus, the contribution of $J_{k \ell}$ to $I E_{r}$ will be less or equal to $u$. In consequence, the more precise $I_{k}$ is, the more precise $I E$ will be.

This shows that it is important to take care of the way the $I_{k}$ and the $J_{k}$ are computed. We will focus on this point in the two following paragraphs.

### 3.2 Interval evaluation of functions - computation of $J_{k}$

Our goal is to compute out of $I_{k}$ an interval $J_{k}$ as precisely as possible with regard to $\varepsilon\left[I_{k}\right]$. We can suppose that $I_{k}$ is a small interval, i.e. in practice its diameter is a lot smaller than 1 . We use the library MPFI ${ }^{\mathbb{1}}$ which implements the interval arithmetic with arbitrary precision. The precision used in the computations is a parameter of our algorithm. For each function $f$

[^2]known by MPFI, and for each interval $I$, the evaluation of $f$ on $I$ by MPFI produces an interval $J$ such that $f[I] \subseteq J$. We will denote by $f(I)$ the interval computed by MPFI.

MPFI implements the standard functions $+,-, /, \times \sqrt{ } \cdot, \exp$, sin, etc. For more complex functions such as $h(x)=\exp (\sin (x+\ln (x)))$ we have to decompose the expression and evaluate each subterms separately. For example, to compute $h$ on $I$, we will first compute $J_{1}=\ln (I)$, then $J_{2}=I+J_{1}$, then $J_{3}=\sin \left(J_{2}\right)$ and finally $J_{4}=\exp \left(J_{3}\right)$. For each operation, MPFI is very precise. In contrast, the interval evaluation of composed functions is subject to error accumulation and cancellation effects caused by decorrelation. The final result of such an evaluation may thus be inaccurate. Shortly speaking, $\operatorname{diam}(J)=\mathcal{O}(\operatorname{diam}(I))$.

We can use the mean value theorem for obtaining an evaluation satisfying $\operatorname{diam}(J)=$ $\mathcal{O}\left(\operatorname{diam}(I)^{2}\right)$, which performs well for the used small $I_{k}$. Let $z$ be a point in $I_{k}$ (we choose the middle of $I_{k}$ ), for each $x \in I_{k}, \exists c \in I_{k}, \varepsilon(x)=\varepsilon(z)+(x-z) \varepsilon^{\prime}(c)$. It follows that $\varepsilon\left[I_{k}\right] \subseteq \varepsilon(z)+\left(I_{k}-z\right) \varepsilon^{\prime}\left[I_{k}\right]$. So we can compute $\varepsilon^{\prime}\left(I_{k}\right)$ using MPFI and then take for $J_{k}$ the interval $\varepsilon(z)+\left(I_{k}-z\right) \varepsilon^{\prime}\left(I_{k}\right)$. The interest of this method comes from the fact that the errors in the evaluation of $\varepsilon^{\prime}\left(I_{k}\right)$ are multiplied by $\left(I_{k}-z\right)$ which is a very small interval centered in 0 .

Obviously, we can use this technique recursively and compute $\varepsilon^{\prime}\left(I_{k}\right)$ using $\varepsilon^{\prime \prime}\left(I_{k}\right)$ and so on. This allows theoretically to obtain $\operatorname{diam}(J)=\mathcal{O}\left(\operatorname{diam}(I)^{n}\right)$. A possible problem is that the successive derivatives of $\varepsilon$ are more and more complex expressions and their evaluation may lead to so unprecise results whilst using great amounts of memory; thus, the technique is useless. At the moment, the number of step of recursivity is just a parameter of the algorithm that the user can fix following its intuition about the complexity of the successive derivatives.

In order to limit the explosion of the expression of the successive derivatives of a function, we implement an additional special optimization for fractions $\varepsilon(x)=\frac{f(x)}{g(x)}$. As long as $g(x)$ has no zero in the given interval, instead of evaluating $\frac{f(z)}{g(z)}+(x-z) \cdot \frac{g(x) \cdot f^{\prime}(x)-f(x) \cdot g^{\prime}(x)}{g^{2}(x)}$, we evaluate $\frac{f(z)+(x-z) \cdot f^{\prime}(x)}{g(z)+(x-z) \cdot g^{\prime}(x)}$. This is more performant since the induced expression trees are smaller than the tree for $\frac{g(x) \cdot f^{\prime}(x)-f(x) \cdot g^{\prime}(x)}{g^{2}(x)}$.

Another problem can arise: some functions have a so-called removable singularity: at some point $z$, the function $\varepsilon(x)$ is of the form $\frac{f(z)}{g(z)}$ with $f(z)=g(z)=0$. However, the function may be prolongated by continuity. It is the case, for instance, for the function $\frac{\sin (x)}{x}$ at 0 . Mathematically, the function remains well defined, but numerically, will perform very badly. The errors of computation become very big; if using interval arithmetic, a division by an interval containing 0 occurs and produces a NaN or an infinity. In order to solve the problem, we have to detect this case and find a solution. If we can detect it (that is if we find a point $z$ and we can prove that $f(z)=g(z)=0$ ), we can use a variant of the so-called L'Hôpital's rule:

$$
\begin{gathered}
\forall x \in I, \exists(c, d) \in I^{2}, \\
\frac{f(x)}{g(x)}=\frac{f(z)+(x-z) f^{\prime}(c)}{g(z)+(x-z) g^{\prime}(d)}=\frac{f^{\prime}(c)}{g^{\prime}(d)} .
\end{gathered}
$$

Thus $(f / g)[I] \subseteq\left(f^{\prime} / g^{\prime}\right)[I]$. Once again, we can use the rule recursively if $f^{\prime} / g^{\prime}$ has a removable singularity in the interval.

For detecting a removable singularity, we firstly test if the function to evaluate is a quotient. If so, we evaluate an interval $J$ containing the denominator $g[I]$ (using Taylor and MPFI). If $J$ does not contain 0 , we are sure that there is no singularity. If it contains 0 , there is a
doubt: we search a floating-point (not interval) zero using the Newton-Raphson method. If we do not find any, we cannot do anything. But if we find one $z$, it is a potential removable singularity. We first use MPFI to evaluate $g$ on the interval $[z, z]$. If the result is $[0,0]$ we know that $z$ is a zero of $g$ (if it is not, we cannot conclude). Since we are now sure that $z$ is a zero of $g$, we evaluate $f$ on $[z, z]$ with MPFI and if the result is $[0,0]$, we know that $z$ is a removable singularity and L'Hôpital's rule can be applied. In every other case, we just let the normal interval evaluation continue, leading to a final result that is NaN or infinity.

It can be argued that this technique is useless because its detection is subject to too much floating-point noise. If it works, it would just be luck. We must have simultaneously discovered by a floating-point Newton-Raphson the precise real point $z$, obtain $f([z, z])=[0,0]$ and $g([z, z])=[0,0]$ using MPFI. This is right but it is the only way to be sure that we actually have a removable singularity. Besides, let us recall that we need to be sure in order to prove the correctness of our final result. Moreover, it works more often that it seems. Actually, if the singularity is not at a floating-point number, the only way to show the infinite norm is finite is to show formally that $g(x)=0$. Further, the function would be so ill-conditioned near the singularity, that it cannot be evaluated in practice. So, Newton-Raphson will almost surely detect the point for functions that are practically evaluated e.g. in elementary function libraries. In addition, $z$ will probably be some simple point such as 0 or an integer and if the functions are not to complicated, it is probable that for this special point, MPFI will be infinitely precise during all the computation. For example,

$$
\frac{\exp (\arcsin (x))-1}{\arcsin (x)}
$$

works well because 0 is detected as a potential singularity, and MPFI knows that $\arcsin ([0,0])=$ $[0,0], \exp ([0,0])=[1,1]$, etc.

The algorithm for evaluating a composed function $\varepsilon$ on an interval $I$ for a result $J$ satisfying $\operatorname{diam}(J)=\mathcal{O}\left(\operatorname{diam}(I)^{r+1}\right.$.
evaluate: input an expression representing a function $\varepsilon$, an interval $I$, and a parameter rec_level:

1. if rec_level>0: differentiate $\varepsilon$; compute the mid-point $z$ of the interval $I$. Returns evaluate $(\varepsilon,[z, z], 0)+(I-z)$. evaluate $\left(\varepsilon^{\prime}, I\right.$, rec_level-1) using MPFI for the addition and the multiplication.
2. else:
(a) if $\varepsilon$ is not a quotient: $\varepsilon$ is of the form $\operatorname{op}(h)$ (or $h_{1}$ op $h_{2}$ ). Return op(evaluate ( $h, I, 0$ )) performing op with MPFI (idem if there is two operands).
(b) else $\varepsilon=h_{1} / h_{2}$. Let $J_{2}=\operatorname{evaluate}\left(h_{2}, I, 0\right)$.
i. If $J_{2}$ does not contain 0 , let $J_{1}=$ evaluate $\left(h_{2}, I, 0\right)$ and return $J_{1} / J_{2}$ performed by MPFI.
ii. else test if L'Hôpital's rule can be applied as explained above. If the test succeeds, formally differentiate $h_{1}$ and $h_{2}$ and return evaluate $\left(h_{1}^{\prime} / h_{2}^{\prime}, I, 0\right)$.

The actual implementation of this algorithm integrates some additional improvements with regard to the performance and the accuracy of the produced results:

- Formally differentiated expressions are simplified exactly as well as possible. The simplification comprises the evaluation of constant sub-expressions as long as no rounding occurs, conversion of polynomial sub-expressions into Horner's scheme, elimination of additions and subtractions with 0 , multiplications and divisions with 0 and 1 and some other simple formal simplifications.
- If their necessity is known in advance, functions are differentiated only once for several evaluations.
- Cancellation in additions and subtractions during interval evaluation of sub-expressions are detected if possible by simple tests. Intermediate Taylor evaluations allow here for improving the accuracy of the result.


### 3.3 Intervals bounding the zeros of a function - determination of $I_{k}$

We have seen that the intervals $I_{k}$ should be small in order to contribute efficiently to the inner enclosure. For the economy of useless computations, we should try to select only those intervals which actually contain some zeros of the derivative. In contrast, for ensuring the correctness of the algorithm, we have to be sure that every zero lies in a $I_{k}$.

In our approach we therfore fix an appropriate diameter $\delta$ as a parameter of the algorithm. We use a dichotomic algorithm. At first, we evaluate $\varepsilon^{\prime}$ on the whole interval $I$ (using all the optimisations of the evaluate function). If the returned interval $J$ contains 0 , we cut $I$ in two parts $I_{1}$ and $I_{2}$ and we recurse on each sub-interval.

If we find an interval $I^{\prime}$ such that $J^{\prime}$ does not contain 0 at one moment of this procedure, then we are sure that $\varepsilon^{\prime}$ has no root in $I^{\prime}$ and we can just eliminate the interval $I^{\prime}$. We stop branching when we have an interval $I^{\prime}$ which diameter is less than $\delta$.

At the end of the procedure, we get a list $\mathcal{I}_{1}, \cdots, \mathcal{I}_{t}$ such that every zero of $\varepsilon^{\prime}$ lies in one $\mathcal{I}_{k}$ and which diameters are all less than $\delta$.

Let $z$ be a zero of $\varepsilon^{\prime}$ and $\mathcal{I}_{k}=[a, b]$ the selected interval in which it lies. In practice, the algorithm computes very often additional intervals $\mathcal{I}_{k-1}$ and $\mathcal{I}_{k+1}$ that actually do not contain any zero of $\varepsilon^{\prime}$ and are of the form $\left[a^{\prime}, a\right]$ and $\left[b, b^{\prime}\right]$ because $\varepsilon\left[\mathcal{I}_{k-1}\right]$ and $\varepsilon\left[\mathcal{I}_{k+1}\right]$ are too close to zero to be discarded by interval evaluation. In contrast, if $\varepsilon$ has a removable singularity in $z$, the evaluation of $\varepsilon\left(\mathcal{I}_{k-1}\right)$ and $\varepsilon\left(\mathcal{I}_{k+1}\right)$ will be very unstable since $\varepsilon$ is ill-conditioned near $z$ and may yield to unprecise results. However, if we join the intervals, obtaining one interval $I^{\prime}=\mathcal{I}_{k-1} \cup \mathcal{I}_{k} \cup \mathcal{I}_{k+1}$, we can apply L'Hôpital's rule on the whole interval $I^{\prime}$ which yields to better results, even if the interval is three times wider than the previous one.

Thus, we replace every series of consecutive intervals in the list $\mathcal{I}_{1}, \cdots, \mathcal{I}_{t}$ by their union unless the union becomes more than 4 times greater in diameter than the parameter $\delta$ fixed previously. This yields to the final list $I_{1}, \cdots, I_{p}$ used in the following of the algorithm.

Remark that instead of using a dichotomy to bracket the zeros of the derivative $\varepsilon^{\prime}(x)$ of the given function, the interval Newton method as described in [5] could also be used. We would, nevertheless, have to require the input functions $\varepsilon$ to be at least $C^{3}$ in this case.

## 4 Generating a proof for infinite norm results

Interval arithmetic, satisfying the so-called inclusion property, has strong links to numerical proving of mathematical properties. As shown in [3], libraries for certifying interval computations in proof checkers, such as for example the Prototype Verification System (PVS) [9] exist.

The general idea consists in retaining the complete computational and decisional tree of an instance of an interval algorithm, for instance, our infinite norm algorithm, and to generate a lemma for each invokation of a interval function or logical element such as dichotomy, Taylor series expansion, L'Hôpital's rule etc.

Our algorithm is currently not yet capable of producing PVS or COQ [7] readable proofs. Nevertheless, it is already possible to store the complete computational tree and to generate an English written proof for an instance of the infinite norm algorithm. The last element that is still lacking to us to provide this additional safety to the correctness of the results computed and inherently proven by the interval algorithm is the difficulty to handle transcendental functions in formal proof checkers. Such a library is partially available for PVS but as shown in [3], computation times for proof checking are still very high.

Most current proofs are still untracktable in PVS because of the complexity of the implied numbers. We are working to provide a means of simplifying a proof in terms of the bitlength of used numbers.

In some examples, the proof generated by our algorithm could already be checked. This is due to the fact that the derivatives of some transcendental functions such as $\log (x)$ are rational.

## 5 Examples

Let us give some examples of the behaviour of our algorithm. The examples are mainly taken out of problems in the development of the crlibm library for correctly rounding elementary functions[1].

1. The first example is a toy problem: let be $\varepsilon(x)=\frac{\log (1+x)}{x}$. The function $\varepsilon(x)$ is defined in 0 only by continuous extension. Our algorithm answers for $\|\varepsilon(x)\|_{\infty}^{\left[-2^{-6} ; 2^{-6}\right]}$ the following

$$
\|\varepsilon(x)\|_{\infty}^{\left[-2^{-6} ; 2^{-6}\right]} \in\left[541109425 \cdot 2^{-29} ; 270554713 \cdot 2^{-28}\right]
$$

This is equivalent to an accuracy of 29 correct bits; the computing precision has been 30 bits. Computation time is less than 1 second on current desktop machines.
2. The second example is the computation of a bound for the approximation error of the polynomial $p(x)=x-\frac{1}{2} \cdot x^{2}+6004799503160663 \cdot 2^{-54} \cdot x^{3}-9007199254173073 \cdot 2^{-55} \cdot x^{4}+$ $3602879701310655 \cdot 2^{-54} \cdot x^{5}-6004904200786859 \cdot 2^{-55} \cdot x^{6}+40211673751819 \cdot 2^{-48} \cdot x^{7}$ with respect to the function $f(x)=\log (1+x)$ in the domain $I=\left[-129 \cdot 2^{-15} ; 129 \cdot 2^{-15}\right]$. Our algorithm returns a lower and upper bound which are close enough that they can both be considered as a result accurate up to 33 bits. The computing precision has been 500 bits. Remark here that $\varepsilon(x)$ is defined in 0 only by continuous extension and ill-conditioned around this point. Computational time is around 3 minutes on current desktop machines.
3. The third and last example shows that results returned by Maple's infinite norm may have no correct bit. Let be $p(x)=x-9223372036854776725 \cdot 2^{-64} \cdot x^{2}+6148914691236520117$. $2^{-64} \cdot x^{3}-18446744071800930591 \cdot 2^{-66} \cdot x^{4}+7378697627908458209 \cdot 2^{-65} \cdot x^{5}-3074519401226530361 \cdot$ $2^{-64} \cdot x^{6}+5270640148006219133 \cdot 2^{-65} \cdot x^{7}, f(x)=\log (1+x)$ and $\varepsilon(x)=p(x)-f(x)$. Our algorithm returns for $\|\varepsilon(x)\|_{\infty}^{\left[-129 \cdot 2^{-15} ; 129 \cdot 2^{-15}\right]}$ as an upper bound the value $0.13178021 \ldots$.
$10^{-21}$, which is an approximation up to 99 bits according to the returned lower bound. Maple 9.5's infnorm returns a value 4 times higher than our upper bound. Thus, the exponents are not the same and no bit of Maple's result is correct. The intermediate computing precision of Maple and of our algorithm are nevertless adjusted to the same, i.e. approximately 165 bits.

## 6 Conclusion

We have given an algorithm for computing a self-validating interval result for the infinite norm $\|\varepsilon(x)\|_{\infty}^{I}$ of an univariate function $\varepsilon(x) \in C^{2}$ on some domain $I$. The algorithm can handle ill-conditioned functions by overcoming ill-conditioning and resulting high deconcellation by the usage of high intermediate (multi-) precision and (recursive) interval Taylor evaluation. Functions with discontinuities in floating-point values that can be extended by continuity can be handled by the use of L'Hôpital's rule, too. The algorithm proves in this case automatically that the L'Hôpital's rule can be applied. Such functions are common as approximation errors in the development of numerical algorithms, in particular, elementary functions.

The implementation of our algorithm is sufficiently performant for common problems on current machines. Examples taken out of the development of crlibm[1], an implementation of correct rounding elementary functions in double precision, can all be handled in some minutes of computation.

The algorithm can retain its computational and decision tree for the generation of an English written proof of an instance. Such a proof may be used in the certification process of a numerical algorithm analysed with our infinite norm. Currently, generation of a PVS or COQ readable proof is impossible because of the difficulty to handle transcendental functions.

We know some examples of functions, as for instance $\frac{\log \sqrt{x}}{\sqrt{x}}$ whose derivatives, as they are provided by our automatic differentiation algorithm, are numerical unstable and cannot be handled by our algorithm. Nevertheless, these derivatives could be formally simplified and brought to a evaluable form. In future our work, we will try to integrate a little more formal simplification into the presented algorithm.

We are currently lacking knowledge of other algorithms and approaches with similar specifications to compare our algorithm and implementation with. Naturally, this might also be due to the fact that our infinite norm problems arise only in particular situations and that no concurrent approaches exist.

## References

[1] CR-Libm, a library of correctly rounded elementary functions in double-precision. http://lipforge.ens-lyon.fr/www/crlibm/.
[2] ANSI/IEEE. Standard 754-1985 for binary floating-point arithmetic, 1985.
[3] Marc Daumas, Guillaume Melquiond, and César Muñoz. Guaranteed proofs using interval arithmetic. In Paolo Montuschi and Eric Schwarz, editors, Proceedings of the 17th IEEE Symposium on Computer Arithmetic, pages 188-195, Cape Cod, Massachusetts, USA, 2005.
[4] F. de Dinechin, Ch. Q. Lauter, and G. Melquiond. Assisted verification of elementary functions using Gappa. In P. Langlois and S. Rump, editors, Proceedings of the 21st Annual ACM Symposium on Applied Computing - MCMS Track, April 2006.
[5] Eldon Hansen and G. William Walster. Global Optimization Using Interval Analysis, Second Edition, Revised and Expanded. Marcel Dekker, Inc., New York, Basel, 2004.
[6] Nicholas J. Higham. Accuracy and Stability of Numerical Algorithms. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, second edition, 2002.
[7] Gérard Huet, Gilles Kahn, and Christine Paulin-Mohring. The Coq proof assistant: a tutorial: version 8.0, 2004.
[8] J.-M. Muller. Elementary Functions, Algorithms and Implementation. Birkhauser, Boston, 1997.
[9] S. Owre, J. M. Rushby, and N. Shankar. PVS: A prototype verification system. In D. Kapur, editor, 11th International Conference on Automated Deduction (CADE), pages 748-752, Saratoga, NY, June 1992.
[10] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. Numerical Recipes in C, The Art of Scientific Computing, 2nd edition. Cambridge University Press, 1992.
[11] Arnaud Tisserand. Hardware operator for simultaneous sine and cosine evaluation. In Proc. International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 992-995, Toulouse, France, May 2006.


[^0]:    *available at http://lipforge.ens-lyon.fr/www/gappa/
    ${ }^{\dagger}$ cf. www.maplesoft.com
    ${ }^{\ddagger}$ cf. www.mathworks.com

[^1]:    ${ }^{\text {§ }}$ available under the GPL at http://lipforge.ens-lyon.fr/projects/arenaireplot

[^2]:    §distributed under the LGPL at http://gforge.inria.fr/projects/mpfi/

