ON THE EXPONENTIATION OF INTERVAL MATRICES

ALEXANDRE GOLDSZTEJN* AND ARNOLD NEUMAIER†

Abstract. The numerical computation of the exponentiation of a real matrix has been studied intensively. The main objective of a good numerical method is to deal with round-off errors and computational cost. The situation is more complicated when dealing with interval matrix exponentiation: Indeed, the main problem will now be the dependence loss in the interval evaluation due to multiple occurrences of some variables, which may lead to enclosures that are too wide to be useful. In this paper, the problem of computing a sharp enclosure of the interval matrix exponential is proved to be NP-hard. Then the scaling and squaring method is adapted to interval matrices and shown to drastically reduce the dependence loss compared with the interval evaluation of the Taylor series.

Although most of what is presented in this paper seems to be known to the experts, one can find nowhere a coherent source for the results. The present paper fills the gap, and adds numerical examples and new insights.

1. Introduction. The exponentiation of a real matrix allows one to solve initial value problems (IVPs) for linear ordinary differential equations (ODEs) with constant coefficients: given $A \in \mathbb{R}^{n\times n}$, the solution of the IVP defined by $y'(t) = A\ y(t)$ and $y(0) = y_0$ is $y(t) = \exp(tA)\ y_0$, where for any $M \in \mathbb{R}^{n\times n}$,

$$\exp(M) := \sum_{k=0}^{\infty} \frac{M^k}{k!}. \tag{1.1}$$

Linear ODE can be found in many contexts. The numerical computation of the matrix exponential has been intensively studied (see, e.g., [25] ,[4], [6], [16], [9], [1] and references therein). While an approximate computation of (1.1) leads to an approximate solution for the underliying IVP, interval analysis (see Section 2) offers a more rigorous framework: In most practical situations the parameters that define the linear ODE are known with some uncertainty only. In this situation, one usually ends up with an interval of matrices $\boldsymbol{A} = [\underline{A}, \overline{A}] := \{A \in \mathbb{R}^{n\times n} : \underline{A} \le A \le \overline{A}\}$ inside which the actual matrix A is known to be. Then the rigorous enclosure of the solution will be obtained by computing an interval matrix that encloses the exponentiation of the interval matrix $\boldsymbol{A}$.

This leads to the following definition of the exponentiation of an interval matrix:

$$\exp(\boldsymbol{A}) := \square\{\exp(A) : A \in \boldsymbol{A}\}, \tag{1.2}$$

where $\square$ denotes the interval hull, i.e. the smallest interval matrix enclosing this set of real matrices.

The most obvious way of obtaining an interval enclosure of $\exp(A)$ is to evaluate the truncated Taylor series using interval arithmetic and to bound the remainder (cf. Subsection 4.1 for details). However, even with high enough order for the expansion, so that the influence of the remainder is insignificant, the interval evaluation of the Taylor series computes very crude bounds on the exponential of an interval matrix. The reason for this bad behavior of the Taylor series interval evaluation is the dependency loss in the interval evaluation that occurs due to multiple occurrences of variables. In general, one cannot expect to compute (1.2) with good precision: The NP-hardness of this problem is proved in Section 3.

*CNRS, France (Alexandre.Goldsztejn@univ-nantes.fr).
†Faculty of Mathematics, University of Vienna, Austria (Arnold.Neumaier@univie.ac.at).

Two well known techniques can help decreasing the pessimism of the interval evaluation: First, for narrow enough interval inputs, centered forms can give rise to sharper enclosures than the natural interval evaluation. Such a centered form for the matrix exponential was proposed in [22, 23, 24]. However, this centered evaluation dedicated to the interval matrix exponentiation is quite complex and very difficult to follow or implement. Furthermore, there is an error in the proof of Proposition 10M of [22][1] and some unjustified assumptions in Section VII [23].

The second technique consists of formally rewriting the expression so as to obtain a formula more suited to interval evaluation (usually decreasing the number of occurrences of variables). For example, the evaluation of a polynomial in its Horner form is known to improve its interval evaluation [5]. It can be naturally applied to the Taylor series of the matrix exponential and was a ctually used in [23] to exponentiate the center matrix as required in the centered form. We give in Subsection 4.2 a proof of the correctness of the matrix exponential Taylor series Horner evaluation with rigorous bound on the truncation that is much simpler than the one given in [23]. In Subsection 4.3, we extend the well known scaling and squaring process (which consists of rewriting the Taylor series using the formula $\exp M = (\exp M/2^L)^{2^L}$) to the exponentiation of interval of matrices. In addition of the usual benefits of this process, its use in conjunction with interval analysis allows an automatic control of the rounding errors. Furthermore, it is shown to drastically improve the dependence loss in the interval evaluation, hence provides much more accurate and less expensive computations. As explained in Section 5 dedicated to experiments, the enclosure formula based on the scaling and squaring process is not only much simpler than the centered evaluation proposed in [23] but it also provides sharper enclosures.

Notation. Standard notation for interval analysis is used.

In particular, the sets of interval, interval vectors and interval matrices are respectively denoted by $\mathbb{IR}$, $\mathbb{IR}^n$ and $\mathbb{IR}^{n\times n}$. Interval objects are denoted by boldface characters, e.g., $\boldsymbol{x} \in \mathbb{IR}$ or $\boldsymbol{A} \in \mathbb{IR}^{n\times m}$.

2. Interval Analysis. Interval analysis (IA) is a modern branch of numerical analysis that was born in the 1960's. It consists of computing with intervals of reals instead of reals, providing a framework for handling uncertainties and verified computations (see [2, 10, 17, 18, 21] and [11] for a survey). The reader is assumed to be familiar with the basics of interval analysis.

When an expression contains several occurrences of some variables its interval evaluation often gives rise to a pessimistic enclosure of the range. For example, the evaluation of $x + xy$ for the arguments $\boldsymbol{x} = [0, 1]$ and $\boldsymbol{y} = [-1, 0]$ gives rise to the enclosure $[-1, 1]$ of $\{x + xy : x \in \boldsymbol{x}, y \in \boldsymbol{y}\}$ while the evaluation of $x(1 + y)$ for the same interval arguments gives rise to the better enclosure $[0, 1]$ of the same range (the latter enclosure being optimal since the expression $x(1 + y)$ contains only one occurrence of each variables). This overestimation is the consequence of the loss of correlation between different occurrences of the same variables when the expression is evaluated for interval arguments.

In the following, except when explicitly mentioned, the notation $\boldsymbol{A}^k$ denotes $\boldsymbol{A}\,\boldsymbol{A}\cdots\boldsymbol{A}$, which is an enclosure of $\{A^k : A \in \boldsymbol{A}\}$. However, let us note that while $\boldsymbol{A}\,\boldsymbol{B} = \square\{AB : A \in \boldsymbol{A}, B \in \boldsymbol{B}\}$, the interval evaluation of $\boldsymbol{A}\,\boldsymbol{A}$, which encloses

[1] The proof of this Proposition is claimed to be similar to the proof of Proposition 10 of [22]. However, the proof of Proposition 10 uses the fact that $f(\boldsymbol{x}) = \cup_{x\in\boldsymbol{x}} f(x)$, which is valid only for scalar functions but not for vector-valued or matrix-valued functions, and thus cannot be extended to prove Proposition 10M, which involves matrix-valued functions.

$\{A^2 : A \in \boldsymbol{A}\}$, is not optimal in general since several occurrences of some entries of $\boldsymbol{A}$ appear in each expression of the entries of $\boldsymbol{A}\,\boldsymbol{A}$. An algorithm for the computation of $\square\{A^2 : A \in \boldsymbol{A}\}$, which can be evaluated with a number of interval operations that is polynomial w.r.t. to the dimension of $\boldsymbol{A}$, was proposed in [13]. However, it was proved in [13] that no such polynomial algorithm exists for the computation of $\square\{A^3 : A \in \boldsymbol{A}\}$ unless P=NP, i.e., the computation of $\square\{A^3 : A \in \boldsymbol{A}\}$ is NP-hard.

The situation is even worse than this: Even computing an enclosure of $\{A^3 : A \in \boldsymbol{A}\}$ for a fixed precision is NP-hard. The notion of ϵ-accuracy of an enclosure is introduced to formalize this problem (see e.g. [7, 14]). The following definition is adapted to sets of matrices.

Definition 2.1. *Let $\mathbb{A} \subseteq \mathbb{R}^{n\times m}$ be a bounded set of matrices, $\boldsymbol{A} = \square\mathbb{A} \in \mathbb{IR}^{n\times m}$, and consider an interval enclosure $\boldsymbol{B}$ of $\mathbb{A}$ (which obviously satisfies $\boldsymbol{B} \supseteq \boldsymbol{A}$). The interval enclosure $\boldsymbol{B}$ is said ϵ-accurate if*

$$\max\{\ \max_{ij} |\underline{a}_{ij} - \underline{b}_{ij}|\ ,\ \max_{ij} |\overline{a}_{ij} - \overline{b}_{ij}|\ \}\ \le\ \epsilon \tag{2.1}$$

Thus an ϵ-accuracy enclosure of a set of matrices is ϵ-accurate for each entry. Although it is not stated in [13], the proof presented there also shows that the computation of an ϵ-accurate enclosure of $\{A^3 : A \in \boldsymbol{A}\}$ is NP-hard.

3. Computational Complexity of an ϵ-Accurate Interval Matrix Exponential. Computing ϵ-accurate interval enclosures of the range of a multivariate polynomial $f : \mathbb{R}^n \longrightarrow \mathbb{R}$ is NP-hard (cf. [7] and Theorem 3.1 in [14]). Even if one restricts one's attention to bilinear functions, the computation of ϵ-accurate enclosures of their range remains NP-hard (cf. Theorem 5.5 in [14]). Note that if one fixes the dimension of the problems, then the computation of these ϵ-accurate enclosures is not NP-hard anymore, showing that the NP-hardness is linked to the growth of the problem dimension.

Our next result, that computing an ϵ-accurate enclosure of the interval matrix exponential is NP-hard, is therefore not a surprise.

Theorem 3.1. *For every $\epsilon > 0$, computing an ϵ-accurate enclosure of* $\exp(\boldsymbol{A})$ *is NP-hard.*

Proof. Let $B \in \mathbb{R}^{n\times n}$ and $x, y \in \mathbb{R}^n$ and define $A \in \mathbb{R}^{(2n+2)\times(2n+2)}$ by

$$A := \begin{pmatrix} 0 & x^T & 0 & 0 \\ 0 & 0 & 6B & 0 \\ 0 & 0 & 0 & y \\ 0 & 0 & 0 & 0 \end{pmatrix} \tag{3.1}$$

We check that $A^4 = 0$:

$$A^2 = \begin{pmatrix} 0 & 0 & 6x^TB & 0 \\ 0 & 0 & 0 & 6By \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad A^3 = \begin{pmatrix} 0 & 0 & 0 & 6x^TBy \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad A^4 = 0. \tag{3.2}$$

Therefore we have $\exp(A) = I + A + \frac{A^2}{2} + \frac{1}{6}A^3$:

$$\exp(A) = \begin{pmatrix} 1 & x^T & 3x^TB & x^TBy \\ 0 & I & 6B & 3By \\ 0 & 0 & I & y \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{3.3}$$

and the entry $(1, 2n+2)$ of $\exp(A)$ is $x^T\ B\ y$.

Now, define $\boldsymbol{A} = [\underline{A}, \overline{A}] \in \mathbb{IR}^{(2n+2)\times(2n+2)}$ by

$$\underline{A} := \begin{pmatrix} 0 & \underline{x}^T & 0 & 0 \\ 0 & 0 & 6B & 0 \\ 0 & 0 & 0 & \underline{y} \\ 0 & 0 & 0 & 0 \end{pmatrix} \text{ and } \overline{A} := \begin{pmatrix} 0 & \overline{x}^T & 0 & 0 \\ 0 & 0 & 6B & 0 \\ 0 & 0 & 0 & \overline{y} \\ 0 & 0 & 0 & 0 \end{pmatrix}, \tag{3.4}$$

which are obviously built in polynomial time from B. Matrices that belong to $\boldsymbol{A}$ are of the form (3.1); therefore

$$\left\{\left(\exp(A)\right)_{1,2n+2} : A \in \boldsymbol{A}\right\} = \{x^T\ B\ y : x \in \boldsymbol{x}, y \in \boldsymbol{y}\}. \tag{3.5}$$

Thus, the entry $(1, 2n+2)$ of an ϵ-accurate enclosure of $\exp(\boldsymbol{A})$ is an ϵ-accurate enclosure of the image of $\boldsymbol{x}$ and $\boldsymbol{y}$ by the function $x^T\ B\ y$. This proves that the ϵ-accurate enclosure of the range of a bilinear function $f(x, y) = x^T\ B\ y$, which is NP-hard by Theorem 5.5 in [14], reduces to the ϵ-accurate enclosure of the interval matrix exponential. Note that no approximation is involved in the argument itself. □

Roughly speaking, Theorem 3.1 states that computing sharp enclosure of the exponentiation of large interval matrices that have no special structure is not tractable.

4. Polynomial Time Algorithms for the Enclosure of the Interval Matrix Exponential. This section presents three expressions dedicated to the enclosure of the exponential of an interval matrix: The naive interval evaluation of the Taylor series, the interval evaluation of the Taylor series following the Horner scheme, and the interval evaluation of the series following the scaling and squaring process. The following example will illustrate each expression.

EXAMPLE 1. *Consider the interval of matrices* $\boldsymbol{A} := [\underline{A}, \overline{A}]$ *where*

$$\underline{A} := \begin{pmatrix} 0 & 1 \\ 0 & -3 \end{pmatrix} \quad , \quad \overline{A} := \begin{pmatrix} 0 & 1 \\ 0 & -2 \end{pmatrix} \quad \text{and} \quad A(t) := \begin{pmatrix} 0 & 1 \\ 0 & t \end{pmatrix}. \tag{4.1}$$

Computing the formal expression of the exponential of the m atrix $A(t)$ *for* $t \in [-3, -2]$*, it can be proved that* $\exp(\boldsymbol{A}) = [\underline{X}, \overline{X}]$ *with*

$$\begin{aligned} \underline{X} &= \begin{pmatrix} 1 & \frac{1}{3}\left(1 - e^{-3}\right) \\ 0 & e^{-3} \end{pmatrix} \approx \begin{pmatrix} 1 & 0.316738 \\ 0 & 0.0497871 \end{pmatrix} \\ \overline{X} &= \begin{pmatrix} 1 & \frac{1}{2}\left(1 - e^{-2}\right) \\ 0 & e^{-2} \end{pmatrix} \approx \begin{pmatrix} 1 & 0.432332 \\ 0 & 0.135335 \end{pmatrix}. \end{aligned} \tag{4.2}$$

4.1. Taylor Series. The naive interval evaluation of the truncated Taylor series for interval matrix exponential becomes an enclosure if we include a rigorous bound on the truncation error. The bound used here is the same as in [23]. Let us define for $K + 2 > \|\boldsymbol{A}\|$

$$\begin{aligned} \tilde{\boldsymbol{\mathcal{T}}}_K(\boldsymbol{A}) &:= I + \boldsymbol{A} + \tfrac{1}{2}\boldsymbol{A}^2 + \ldots + \tfrac{1}{K!}\boldsymbol{A}^K \\ \boldsymbol{\mathcal{T}}_K(\boldsymbol{A}) &:= \tilde{\boldsymbol{\mathcal{T}}}_K(\boldsymbol{A}) + \boldsymbol{\mathcal{R}}_K(\boldsymbol{A}), \end{aligned} \tag{4.3}$$

where the interval remainder $\boldsymbol{\mathcal{R}}_K(\boldsymbol{A})$ is

$$\boldsymbol{\mathcal{R}}_K(\boldsymbol{A}) := \rho(\|\boldsymbol{A}\|, K)\ [-E, E] \quad \text{with} \quad \rho(\alpha, K) = \frac{\alpha^{K+1}}{(K+1)!\left(1 - \frac{\alpha}{K+2}\right)} \tag{4.4}$$

and $E \in \mathbb{R}^{n\times n}$ has all its entries equal to 1. The interval matrix norm used here and in the remainder of the paper is the infinity norm

$$\|A\| = \max_i \sum_k |\boldsymbol{A}_{ik}| = \max_i \sum_k \max(-\underline{A}_{ik}, \overline{A}_{ik})$$

defined in [21], which is the maximum of the infinite norm of all real matrices included in the interval matrix. This norm is inclusion-increasing, i.e.,

$$\boldsymbol{A} \subseteq \boldsymbol{B} \Rightarrow \|\boldsymbol{A}\| \leq \|\boldsymbol{B}\|.$$

The following results enable us to apply the fundamental theorem of interval analysis to expressions that include $\boldsymbol{\mathcal{R}}_K(\,\cdot\,)$.

LEMMA 4.1. *For any fixed positive integer K, the interval matrix operator $\boldsymbol{\mathcal{R}}_K(\,\cdot\,)$ is inclusion-increasing inside $\{\boldsymbol{A} \in \mathbb{IR}^{n\times n} : \|\boldsymbol{A}\| < K+2\}$.*

Proof. Let $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{IR}^{n\times n}$ such that $\|\boldsymbol{B}\| \leq K+2$ and $\boldsymbol{A} \subseteq \boldsymbol{B}$. Then $\|\boldsymbol{A}\| \leq \|\boldsymbol{B}\|$, since the norm is i nclusion-increasing, which implies $\|\boldsymbol{A}\| \leq K+2$. Furthermore as $\rho(\alpha, K)$ is obviously increasing with respect to α, we have $\rho(\|\boldsymbol{A}\|, K) \leq \rho(\|\boldsymbol{B}\|, K)$. Finally, as $[-E, E]$ is centered on the null matrix, we have

$$\rho(\|\boldsymbol{A}\|, K)\ [-E, E] \subseteq \rho(\|\boldsymbol{B}\|, K)\ [-E, E], \tag{4.5}$$

which concludes the proof. ◻

The next lemma is a direct consequence of the well known upper bound on the truncation error of the exponential series.

LEMMA 4.2. *Let $A \in \mathbb{R}^{n\times n}$ and $K \in \mathbb{N}$ such that $K+2 > \|A\|$. Then* $\exp(A) \in \boldsymbol{\mathcal{T}}_K(\boldsymbol{A})$.

Proof. Suppose that $\exp(A) \notin \boldsymbol{\mathcal{T}}_K(\boldsymbol{A})$, i.e., there exist $i, j \in \{1, \dots n\}$ such that

$$\big(\exp(A)\big)_{ij} \notin \Big(\sum_{k=0}^{K} \frac{A^k}{k!}\Big)_{ij} +\ \rho(\|A\|, K)\ [-1, 1]. \tag{4.6}$$

This obviously implies

$$\big|\big(\exp(A)\big)_{ij} - \big(\sum_{k=0}^{K} \frac{A^k}{k!}\big)_{ij}\big| > \rho(\|A\|, K). \tag{4.7}$$

Therefore $\|\exp(A) - \sum_{k=0}^{K} \frac{A^k}{k!}\| > \rho(\|A\|, K)$ holds, which contradicts the well known bound on the truncation error for the exponential series (see e.g. [23]). Therefore $\exp(A) \in \boldsymbol{\mathcal{T}}_K(\boldsymbol{A})$ has to hold. ◻

Theorem 4.3 below states that $\boldsymbol{\mathcal{T}}_K(\boldsymbol{A})$ is an enclosure of $\exp(\boldsymbol{A})$. It was stated in [23] but proved with different arguments in [22], [23]. Note that the usage of the fundamental theorem of interval analysis[2] allows us to provide a proof much simpler than the one given in [22], [23].

THEOREM 4.3. *Let $\boldsymbol{A} \in \mathbb{IR}^{n\times n}$ and $K \in \mathbb{N}$ such that $K+2 > \|\boldsymbol{A}\|$. Then* $\exp(\boldsymbol{A}) \subseteq \boldsymbol{\mathcal{T}}_K(\boldsymbol{A})$.

Proof. First, by Lemma 4.1 $\boldsymbol{\mathcal{R}}_K(\,\cdot\,)$ is inclusion-increasing, and therefore so is $\boldsymbol{\mathcal{T}}_K(\,\cdot\,)$ because it is composed of inclusion-increasing operations. Second, by Lemma

[2]The fundamental theorem of interval analysis is the basic argument stating that if an interval function $\boldsymbol{F}$ is inclusion-increasing and if $F(x) \in \boldsymbol{F}(x)$ for some real function F, then $\boldsymbol{F}(\boldsymbol{x}) \supseteq \{F(x) : x \in \boldsymbol{x}\}$; see, e.g., Theorem 1.4.1 of [21].

4.2, $(\forall A \in \boldsymbol{A})\ \exp(A) \in \boldsymbol{\mathcal{T}}_K(\boldsymbol{A})$. Therefore, one can apply the fundamental theorem of interval analysis to conclude the proof. □

EXAMPLE 2. *Consider the interval of matrices $\boldsymbol{A}$ defined in Example 1. Theorem 4.3 with $K = 16$ gives rise to the following enclosure of* $\exp(\boldsymbol{A})$*:*

$$\begin{pmatrix} 1 + [-9 \times 10^{-7}, 9 \times 10^{-7}] & [-1.2092, 1.9582] \\ [-9 \times 10^{-7}, 9 \times 10^{-7}] & [-6.2557, 6.4409] \end{pmatrix}. \tag{4.8}$$

Using a higher order for the expansion do not improve the entries $(1, 2)$ and $(2, 2)$ anymore.

4.2. Horner scheme. The Horner evaluation of a real polynomial improves both the computation cost and the numerical stability (see, e.g., [12]). When an interval evaluation is computed, the Horner evaluation can furthermore reduce the effect of the loss of correlation (see [5]). It is therefore natural to evaluate (4.3) using a Horner scheme:

$$\begin{aligned} \tilde{\boldsymbol{\mathcal{H}}}_K(\boldsymbol{A}) &:= \; I + \boldsymbol{A}\Big(I + \tfrac{\boldsymbol{A}}{2}\Big(I + \tfrac{\boldsymbol{A}}{3}\Big(\;\; \cdots \;\; \Big(I + \tfrac{\boldsymbol{A}}{K}\Big) \cdots \Big)\Big)\Big) \\ \boldsymbol{\mathcal{H}}_K(\boldsymbol{A}) &:= \; \tilde{\boldsymbol{\mathcal{H}}}_K(\boldsymbol{A}) + \boldsymbol{\mathcal{R}}_K(\boldsymbol{A}). \end{aligned} \tag{4.9}$$

LEMMA 4.4. *Let $A \in \mathbb{R}^{n\times n}$ and $K \in \mathbb{N}$ such that $K + 2 > \|A\|$. Then* $\exp(A) \in \boldsymbol{\mathcal{H}}_K(A)$.

Proof. When interval operations are evaluated with real arguments, the Horner scheme can be expanded exactly, leading to $\boldsymbol{\mathcal{H}}_K(A) = \tilde{\boldsymbol{\mathcal{T}}}_K(A)$. As a consequence, $\boldsymbol{\mathcal{H}}_K(A) = \boldsymbol{\mathcal{T}}_K(A)$ and Lemma 4.2 concludes the proof. □

THEOREM 4.5. *Let $\boldsymbol{A} \in \mathbb{IR}^{n\times n}$ and $K \in \mathbb{N}$ such that $K + 2 > \|\boldsymbol{A}\|$. Then* $\exp(\boldsymbol{A}) \in \boldsymbol{\mathcal{H}}_K(\boldsymbol{A})$.

Proof. As a consequence of Lemma 4.1, $\boldsymbol{\mathcal{H}}_K(\,\cdot\,)$ is composed of inclusion-increasing operations, hence is inclusion increasing. Lemma 4.4 shows that $\exp(A) \in \boldsymbol{\mathcal{H}}_K(A)$. Therefore one can use the fundamental theorem of interval analysis to conclude the proof. □

EXAMPLE 3. *Consider the interval of matrices $\boldsymbol{A}$ defined in Example 1. Theorem 4.5 with $K = 16$. We have the following enclosure of* $\exp(\boldsymbol{A})$*:*

$$\begin{pmatrix} 1 + [-1.1 \times 10^{-6}, 1.1 \times 10^{-6}] & [-0.0706, 0.7352] \\ [-1.1 \times 10^{-6}, 1.1 \times 10^{-6}] & [-1.2056, 1.2117] \end{pmatrix}. \tag{4.10}$$

This enclosure is sharper than the one computed using the Taylor series: as it was foreseen, the Horner evaluation actually reduces the effect of the loss of dependency in the expression of the Taylor expansion of the matrix exponential.

4.3. Scaling and squaring process. The scaling and squaring process is one of the most efficient way to compute a real matrix exponential. It consists of first computing $\exp(A/2^L)$ and then squaring L times the resulting matrix:

$$\exp(A) = \big(\exp(A/2^L)\big)^{2^L}. \tag{4.11}$$

Therefore, one first has to compute $\exp(A/2^L)$. This computation is actually much easier than that of $\exp(A)$ because $\|A/2^L\|$ can be made much smaller than 1. Usually, Padé approximations are used to compute $\exp(A/2^L)$. However, this technique has not been extended to interval matrices, hence we propose here to use instead the Horner evaluation of the Taylor series. More specifically, we propose the following operator for the enclosure of an interval matrix exponential: Let K and L be such that $(K+2)2^L > \|\boldsymbol{A}\|$, and define

$$\boldsymbol{\mathcal{S}}_{L,K}(\boldsymbol{A}) := \big(\boldsymbol{\mathcal{H}}_K(\boldsymbol{A}/2^L)\big)^{2^L}. \tag{4.12}$$

The exponentiation in (4.12) is of course computed with L successive interval matrix square operations (here a simple multiplication of the matrix by itself is used).

THEOREM 4.6. *Let* $\boldsymbol{A} \in \mathbb{IR}^{n\times n}$ *and* $K, L \in \mathbb{N}$ *such that* $(K+2)\,2^L > \|\boldsymbol{A}\|$. *Then* $\exp(\boldsymbol{A}) \subseteq \boldsymbol{\mathcal{S}}_{L,K}(\boldsymbol{A})$.

Proof. By Theorem 4.5, we have $\exp(A/2^L) \in \boldsymbol{\mathcal{H}}_K(\boldsymbol{A}/2^L)$ for an arbitrary $A \in \boldsymbol{A}$. The interval evaluation $\boldsymbol{M}^{2^L}$ of an arbitrary interval matrix $\boldsymbol{M}$ encloses $\{M^{2^L} : M \in \boldsymbol{M}\}$; therefore $\boldsymbol{\mathcal{S}}_{L,K}(\boldsymbol{A}) \ni \exp(A/2^L)^{2^L}$. This holds for an arbitrary $A \in \boldsymbol{A}$, which concludes the proof. □

EXAMPLE 4. *Consider the interval of matrices* $\boldsymbol{A}$ *defined in Example 1. Theorem 4.6 with* $L = 10$ *and* $K = 10$ *leads to the following enclosure of* $\exp(\boldsymbol{A})$*:*

$$\begin{pmatrix} 1 + [-5.7\times 10^{-13}, 9.1\times 10^{-13}] & [0.3165, 0.4325] \\ [-2.4\times 10^{-19}, 2.4\times 10^{-19}] & [0.0496, 0.1355] \end{pmatrix}. \tag{4.13}$$

This enclosure is much sharper than the two previously computed enclosures using the Taylor series (cf. Example 2) and its Horner evaluation (cf. Example 3). It is also very close to the optimal enclosure (4.2). *The computation cost for L and K is approximately the same as that for the Horner scheme with order* $L+K$.

5. Experiments. The direct Taylor series is not presented as it is similar but always worse that its Horner evaluation. Therefore we compare in this section only the interval Horner evaluation of the truncated Taylor series with the interval scaling and squaring method. In order to compare these two enclosures, we use the width of these interval enclosure: Let wid $\boldsymbol{M}$ be the real matrix formed of the widths of the entries of $\boldsymbol{M}$. We will use the $\|\text{wid}\,\boldsymbol{M}\|$ as a quality measure of the enclosure $\boldsymbol{M}$. Subsection 5.1 presents a detailed study of the exponentiation of a particular real matrix of dimension 3, while Subsection 5.2 deals with the exponentiation of an interval matrix of dimension 3. Finally Subsection 5.3 presents the interval exponentiation of tridiagonal interval matrices of dimensions up to 100.

Except when explicitly mentioned, we use the following heuristic choice for the order of the expansions of the Taylor series: $\boldsymbol{\mathcal{H}}(\boldsymbol{A})$ denotes $\boldsymbol{\mathcal{H}}_K(\boldsymbol{A})$ where K is chosen as the smallest integer such that $\rho(\|\boldsymbol{A}\|, K) \le 10^{-16}$. Furthermore, $\boldsymbol{\mathcal{S}}(\boldsymbol{A})$ denotes $\boldsymbol{\mathcal{S}}_{S,K}(\boldsymbol{A})$ where $S = \lceil \log_2(10\|\boldsymbol{A}\|)\rceil$ (so that S is the smallest integer such that $\|\boldsymbol{A}/2^S\| \le 0.1$) and $K = 9$ (applying the previous heuristic with $\|\boldsymbol{A}\| = 0.1$). This choice may lead to some overscaling, but this is not discussed here (see, e.g., [1] and references therein for more elaborated heuristics).

As explained in introduction, the comparisons presented in this section do not include the interval matrix enclosure method proposed in [22, 23, 24]. However, this method is based on an interval Horner evaluation and thus cannot give rise to better enclosures than the interval Horner evaluation of the center matrix, which – as demonstrated above – is of poor quality.

5.1. Interval exponentiation of a real matrix. In this subsection, we consider the matrix A defined by

$$A := \begin{pmatrix} -131 & 19 & 18 \\ -390 & 56 & 54 \\ -387 & 57 & 52 \end{pmatrix} \tag{5.1}$$

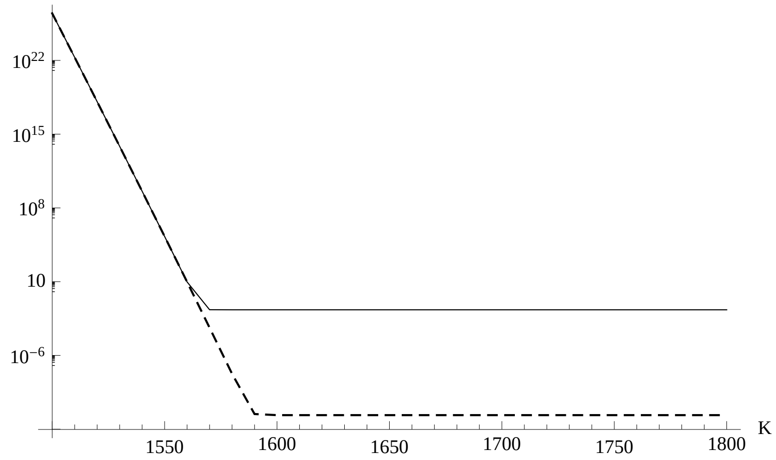

FIG. 5.1. *Plots of* $\|\text{wid}\,\boldsymbol{\mathcal{H}}_K(A)\|$ *w.r.t. to* K*, for two different precisions: Plain line for* $p = 110$ *digits precision, and dashed line for* $p = 120$ *digits precision.*

proposed in [4]. This matrix has a significant eigenvalue separation and a poorly conditioned eigenvector set.

The matrix A is too difficult to exponentiate using the Horner evaluation of the truncated interval Taylor series: As $\|A\| = 500$ the Taylor series requires an expansion of order greater than 502. Computing $\boldsymbol{\mathcal{H}}_K(A)$ using double precision does not provide any meaningful enclosure. Figure 5.1 shows the quality of the enclosure obtained for different orders ranging from 1500 to 1800 and two different precisions for computations (the Mathematica [26] arbitrary precision interval arithmetic was used). It shows that no meaningful enclosure is obtained for precision less than $p = 110$ digits or order less than $K = 1550$ using the Horner interval evaluation of the Taylor expansion.

The interval scaling and squaring formula gives rise to $\|\text{wid}\,\boldsymbol{\mathcal{S}}(A)\| \approx 1.5 \times 10^{-5}$ computed using the standard double precision arithmetic. As noted in [6], preprocessing the matrix exponential using a matrix decomposition can improve the stability of the computation. If one uses a decomposition PMP^{-1} where $M = P^{-1}AP$ is easier to exponentiate then one can compute $\exp A = P \exp(M) P^{-1}$. This gives a rigorous enclosure for arbitrary P, as long as P^{-1} is rigorously enclosed. In particular, we may choose P to be the result of an approximate matrix decomposition. Using the Schur decomposition, we obtain

$$\|\text{wid}\,\big(P\,\boldsymbol{\mathcal{S}}(P^{-1}AP)\,P^{-1}\big)\| \approx 1.1 \times 10^{-10}. \tag{5.2}$$

5.2. Interval exponentiation of an interval matrix. In order to compare the different methods, we will use $0.1A$, where A is given by (5.1), which is simpler to exponentiate than (5.1). We have exponentiated $\boldsymbol{A}_\epsilon := 0.1A + [-\epsilon, \epsilon]$ for various values of ϵ inside $[10^{-16}, 1]$; the results are plotted in Figure 5.2. The three plain gray curves represent $\|\text{wid}\,\boldsymbol{\mathcal{H}}_K(\boldsymbol{A}_\epsilon)\|$ for $K = 150$, $K = 160$ and $K = 170$. Increasing K improves the enclosure until $K = 170$ above which no significantly improvement is found. The black curve represents $\|\text{wid}\,\boldsymbol{\mathcal{S}}(\boldsymbol{A}_\epsilon)\|$. The dashed line represents

$$\|\text{wid}\,\square\{\exp \underline{A}_\epsilon, \exp \overline{A}_\epsilon\}\|, \tag{5.3}$$

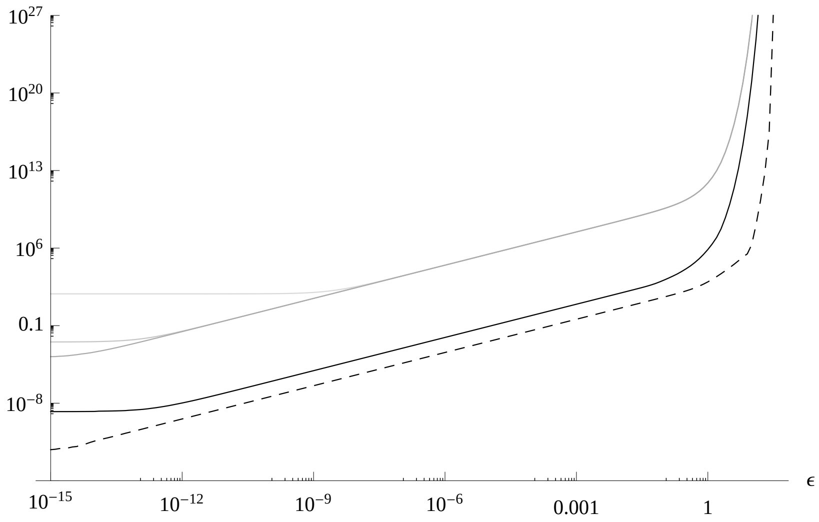

FIG. 5.2. *Norm of the width of the different interval of* $\exp(\boldsymbol{A}_\epsilon)$ *enclosures with respect to* ϵ.

which is a lower bound of $\|\text{wid } \exp(\boldsymbol{A}_\epsilon)\|$. Each plot show three phases: The first phase shows flat plots, then $\|\text{wid}(\,\cdot\,)\|$ increases linearly[3] w.r.t. ϵ before it is eventually exponentially increasing.

During the flat phase, the rounding errors represent the main contribution to the final width of the enclosures. Thus decreasing ϵ does not decrease the width of the final enclosure.

The linear phase is the most interesting. For these values of ϵ, the width of $\exp(\boldsymbol{A}_\epsilon)$ grows linearly because the contribution of quadratic terms are negligible. On the other hand, the interval evaluations $\boldsymbol{\mathcal{H}}_{170}(\boldsymbol{A}_\epsilon)$ and $\boldsymbol{\mathcal{S}}(\boldsymbol{A}_\epsilon)$ are pessimistic. It is well known that the pessimism of interval evaluation grows linearly w.r. to the width of the interval arguments. Indeed, the computed enclosure shows a linear growth w.r. to ϵ which are approximately

$$\|\text{wid}\, \boldsymbol{\mathcal{H}}_{170}(\boldsymbol{A}_\epsilon)\| \approx 1.17 \times 10^{-4} + 2.86 \times 10^{10}\, \epsilon \tag{5.4}$$

$$\|\text{wid}\, \boldsymbol{\mathcal{S}}(\boldsymbol{A}_\epsilon)\| \approx 1.80 \times 10^{-9} + 8.59 \times 10^{3}\, \epsilon. \tag{5.5}$$

This clearly shows how much smaller is the pessimism introduced by the interval scaling and squaring process.

Finally, both the interval Horner evaluation and the interval scaling and squaring process show an exponential growth when ϵ is too large. The lower bound represented by the dashed line also shows an exponential growth, which proves that this is inherent to the exponentiation of an interval matrix. For such ϵ, some matrices inside $\boldsymbol{A}_\epsilon$ eventually see some of their eigenvalues becoming positive, leading to some exponential divergence of the underlying dynamical system, which is also observed in the matrix exponential.

[3] The linear plot displayed within the log-log scale indicates a polynomial behavior, the polynomial degree being fixed by the slope in the log-log representation. Here, the slope is 1 inside the log-log plot and thus so is the degree of the polynomial.

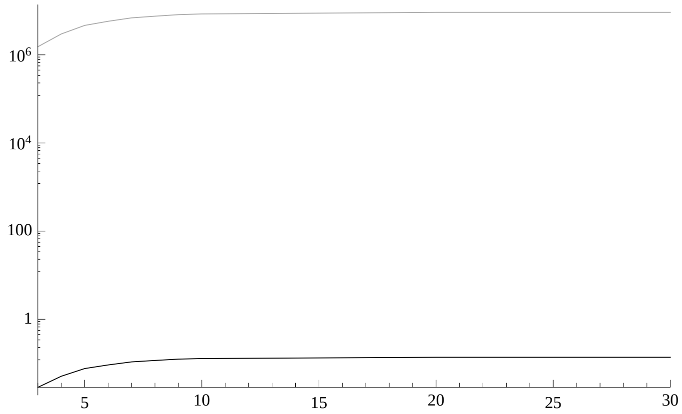

FIG. 5.3. *Norm of the width with respect to n of $\boldsymbol{\mathcal{S}}(\boldsymbol{A}_n)$ (in black) and $\boldsymbol{\mathcal{H}}(\boldsymbol{A}_n)$ (in gray).*

5.3. Tridiagonal Interval Matrices. Consider the tridiagonal interval matrix $\boldsymbol{A}_n \in \mathbb{IR}^{n\times n}$ defined by $t_{ii} = [-11, -9]$ and $t_{i+1,i} = t_{i,i+1} = [0, 2]$. For example,

$$\boldsymbol{A}_3 = \begin{pmatrix} [-11,-9] & [0,2] & 0 \\ [0,2] & [-11,-9] & [0,2] \\ 0 & [0,2] & [-11,-9] \end{pmatrix}. \tag{5.6}$$

Figure 5.3 shows the widths of the enclosures of the exponential of $\boldsymbol{A}_n$ for the different methods presented in the previous section and for values of n ranging from 3 to 100. It clearly shows that the naive interval evaluation and the interval evaluation of the Horner form of the Taylor series are too wide to be useful, while the interval scaling and squaring process gives rise to sharp enclosures.

6. Simulation of Linear ODE. Subsection 6.1 shows how the exponential of interval matrices can help rigorously simulate homogeneous and autonomous linear ODE. Subsection 6.2 shows the link between the interval matrix exponentiation and the application of odd/even reduction proposed by Gambill and Skeel in [8], which allows somehow to generalize the scaling and squaring process to non autonomous linear ODE.

6.1. Homogeneous and Autonomous Linear ODE. We consider the uncertain autonomous linear ODE $x'(t) = Ax(t)$, where A is known to belong to $\boldsymbol{A} = \boldsymbol{A}_3$, the tridiagonal matrix defined in Subsection 5.3. For simplicity, we use a simple simulation algorithm which consists of enclosing the states in a boxes by composing successive time steps using interval analysis (more elaborated algorithms like the one implemented in VNODE-LP [19] more complex enclosing sets like, e.g., parallelotopes [19], zonotopes [15] or nonlinear sets [3]). The initial condition is fixed to belong to $x_0 = (1, \dots, 1) \in \mathbb{R}^n$, and define $\boldsymbol{x}_{k+1} = \boldsymbol{\mathcal{S}}(\boldsymbol{A}_n)\boldsymbol{x}_k$. As $\boldsymbol{\mathcal{S}}(\boldsymbol{A}_n) \supseteq \exp(\boldsymbol{A}_n)$, the state at time $t_k = k$ is obviously contained inside $\boldsymbol{x}_k$. Now, as $\|\boldsymbol{\mathcal{S}}(\boldsymbol{A}_n)\| < 1$ for all $n \leq 100$, the sequence $\boldsymbol{x}_{k+1} = \boldsymbol{x}_k$ is proved to converge to 0. The left diagram in Figure 6.1 shows the norm of $\boldsymbol{x}_k$, which indeed converges to 0.

For comparison, let us consider the well known ODE solver VNODE-LP [19], which is able to enclose rigorously the solution of a nonlinear ODE $x'(t) = f(x(t))$ with uncertain parameters. As most of rigorous ODE solvers based on interval analysis (see [20] and references therein), it is based on the interval evaluation of the Taylor

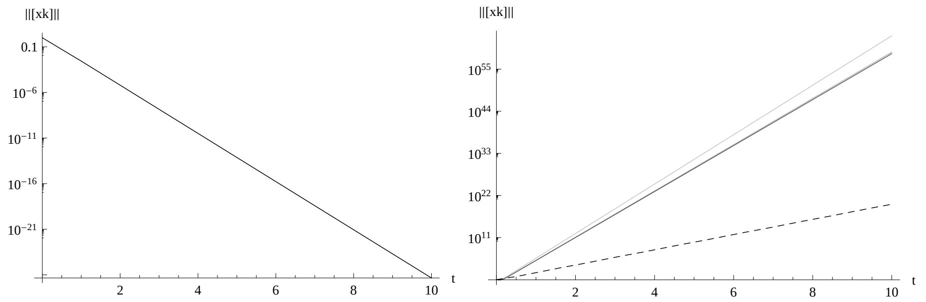

FIG. 6.1. *Norm of the enclosure of a linear ODE solution using various methods: The interval scaling and scaring with time step* 1 *(left in solid line), Equation* (6.4) *(right in solid line using time steps 0.01, 0.001 and 0.0001 from light gray to dark gray) and VNODE-LP (right in dashed line).*

coefficients of the solution: The solution $x(t)$ is enclosed by

$$x_0 + tf^{[1]}(x_0) + t^2 f^{[2]}(x_0) + \cdots + t^m f^{[m]}(x_0) + \boldsymbol{r}, \tag{6.1}$$

where $\boldsymbol{r}$ is an upper bound on the truncation error that can be made arbitrary small for a fixed extension order m using an adaptative stepsize control (which is actually performed by VNODE-LP). The functions $f^{[k]}(x_0)$ correspond to the Taylor coefficients $\frac{1}{k!}x^{(k)}(0)$ expressed w.r.t. x_0. They are evaluated using automatic differentiation of the recursive definition

$$f^{[k]}(x) = \frac{1}{k}\frac{\partial f^{[k-1]}(x)}{\partial x} f(x). \tag{6.2}$$

In the case of a autonomous linear ODE, $f(x) = Ax$ so $f^{[k]}(x) = \frac{1}{k!} A^k$. This leads to the enclosure

$$\boldsymbol{x}_0 + t\boldsymbol{A}\boldsymbol{x}_0 + t^2\big(\frac{\boldsymbol{A}^2}{2}\big)\boldsymbol{x}_0 + \cdots + t^m\big(\frac{\boldsymbol{A}^m}{m!}\big)\boldsymbol{x}_0 + \boldsymbol{r} \tag{6.3}$$

of the state, which is to be compared to the state enclosure obtained using the interval evaluation of the matrix exponential Taylor series, that is

$$\Big(I + t\boldsymbol{A} + \frac{1}{2}(t\boldsymbol{A})^2 + \cdots + \frac{1}{m!}(t\boldsymbol{A})^m\Big)\boldsymbol{x}_0 + \boldsymbol{r}'. \tag{6.4}$$

We may suppose that the time step t is small enough and the order m is large enough so that $\boldsymbol{r}$ and $\boldsymbol{r}'$ are both negligible. Although in the case of real arithmetic both expressions (6.3) and (6.4) are equivalent, the interval evaluation (6.3) is actually much worse than (6.4) in the sense that it leads to a much stronger pessimism. The right hand diagram in Figure 6.1 shows that the simulation indeed diverges when computed using Equation (6.4) for various time steps and using VNODE-LP (which uses an automatic time step selection).

6.2. General Linear ODE. Gambill and Skeel proposed in [8] a method for simulating non-homogeneous and non-autonomous linear ODE that can be interpreted as a generalization of the scaling and squaring process. We summarize their method in the case of homogeneous linear ODE $x'(t) = A(t)x(t)$. For simplicity, let us consider 2^L constant time steps of size h. Gambill and Skeel use (6.1) in order to compute interval matrices $\boldsymbol{G}_{k+1}$ such that $x((k+1)h) \in \boldsymbol{G}_{k+1}x(kh)$. Note that in the case of an autonomous ODE, i.e. $A(t) = A$, $\boldsymbol{G}_k$ is an enclosure of $\exp(hA)$.

They observe that the obvious enclosure of $x(2^L h)$ obtained computing

$$\boldsymbol{G}_{2^L}\Big(\cdots\Big(\boldsymbol{G}_2\Big(\boldsymbol{G}_1 x(0)\Big)\Big)\cdots\Big) \tag{6.5}$$

(called the *marching reduction* in [8]) is useless since it grows exponentially even for very simple problems. The enclosure

$$\boldsymbol{G}_{2^L}\Big(\cdots\Big(\boldsymbol{G}_2\Big(\boldsymbol{G}_1\Big)\Big)\cdots\Big)x(0) \tag{6.6}$$

of $x(2^L h)$ is already much better than (6.5), although it requires more computations since 2^L matrix/matrix products are performed in (6.6) while 2^L matrix/vector products are performed in (6.5). Instead of (6.6), Gambill and Skeel propose to evaluate the successive products starting by multiplying pairs of successive matrices. The operation is repeated L times until the final result is obtained. For example, the enclosure of $x(2^3 h)$ computed by Gambill and Skeel is

$$\Big(\big((\boldsymbol{G}_8\boldsymbol{G}_7)(\boldsymbol{G}_6\boldsymbol{G}_5)\big)\big((\boldsymbol{G}_4\boldsymbol{G}_3)(\boldsymbol{G}_2\boldsymbol{G}_1)\big)\Big)x(0), \tag{6.7}$$

which is much shaper than (6.6) in most situations. It is worthwhile noting that (6.7) reduces to $(\boldsymbol{B})^{2^L}$, where $\boldsymbol{B}$ is a superset of $\exp(hA)$, when the system is autonomous.

Acknowledgments. The first author would like to thank the University of Central Arkansas, USA, who partially funded this work, and in particular Professor Chenyi Hu for his comments.

References.

[1] A. H. Al-Mohy and N. J. Higham. A New Scaling and Squaring Algorithm for the Matrix Exponential. *SIAM J. Matrix Anal. Appl.*, 31(3):970–989, 2009.

[2] G. Alefeld and J. Herzberger. *Introduction to Interval Computations.* Computer Science and Applied Mathematics, 1974.

[3] M. Berz and K. Makino. Verified Integration of ODEs and Flows Using Differential Algebraic Methods on High-Order Taylor Models. *Reliab. Comp.*, 4(4):361–369, 1998.

[4] P. Bochev and S. Markov. A Self-Validating Numerical Method for the Matrix Exponential. *Computing*, 43(1):59–72, 1989.

[5] M. Ceberio and L. Granvilliers. Horner's rule for interval evaluation revisited. *Computing*, 69(1):51–81, 2002.

[6] P. I. Davies and N. J. Higham. A Schur-Parlett Algorithm for Computing Matrix Functions. *SIAM J. Matrix Anal. Appl.*, 25(2):464–485, 2003.

[7] A. A. Gaganov. Computational complexity of the range of the polynomial in several variables. *Cybernetics and Systems Analysis*, 21(4):418–421, 1985.

[8] T. N. Gambill and R. D. Skell. Logarithmic reduction of the wrapping effect with application to ordinary differential equations. *SIAM J. Numer. Anal.*, 25(1):153–162, 1988.

[9] N. J. Higham. The Scaling and Squaring Method for the Matrix Exponential Revisited. *SIAM J. Matrix Anal. Appl.*, 26(4):1179–1193, 2005.

[10] L. Jaulin, M. Kieffer, O. Didrit, and E. Walter. *Applied Interval Analysis with Examples in Parameter and State Estimation, Robust Control and Robotics.* Springer-Verlag, 2001.

[11] R. B. Kearfott. Interval Computations: Introduction, Uses, and Resources. *Euromath*, Bulletin 2(1):95–112, 1996.

[12] D. Knuth. *The Art of Computer Programming, Volume 2: Seminumerical Algorithms (Third Edition).* Addison-Wesley, 1997.

[13] O. Kosheleva, V. Kreinovich, G. Mayer, and H. T. Nguyen. Computing the cube of an interval matrix is NP-hard. In *SAC '05: Proc. of the 2005 ACM Symposium on Applied Computing*, pages 1449–1453, 2005.

[14] V. Kreinovich, A. Lakeyev, J. Rohn, and P. Kahl. *Computational Complexity and Feasibility of Data Processing and Interval Computations.* Kluwer Academic Publishers, 1998.

[15] W. Kühn. Rigorously Computed Orbits of Dynamical Systems Without the Wrapping Effect. *Computing*, 61:47–67, 1998.

[16] C. B. Moler and C. F. Van Loan. Nineteen dubious ways to compute the exponential of a matrix. *SIAM Review*, 45(1):3–49, 2003.

[17] R. E. Moore. *Interval Analysis.* Prentice-Hall, 1966.

[18] R. E. Moore, R. B. Kearfott, and M. J. Cloud. *Introduction to Interval Analysis.* SIAM, 2009.

[19] N. S. Nedialkov. VNODE-LP, A Validated Solver for Initial Value Problems in Ordinary Differential Equations. Technical Report CAS-06-06-NN, McMaster University, 2006.

[20] N. S. Nedialkov, K. R. Jackson, and G. F. Corliss. Validated Solutions of Initial Value Problems for Ordinary Differential Equations. *Applied Mathematics and Computation*, 105(1):21–68, 1999.

[21] A. Neumaier. *Interval Methods for Systems of Equations.* Cambridge Univ. Press, 1990.

[22] E. P. Oppenheimer and A. N. Michel. Application of Interval Analysis Techniques to Linear Systems. I. Fundamental Results. *IEEE Transactions on Circuits and Systems*, 35(9):1129–1138, 1988.

[23] E. P. Oppenheimer and A. N. Michel. Application of Interval Analysis Techniques to Linear Systems. II. The Interval Matrix Exponential Function. *IEEE Transactions on Circuits and Systems*, 35(10):1230–1242, 1988.

[24] E. P. Oppenheimer and A. N. Michel. Application of Interval Analysis Techniques to Linear Systems. III. Initial Value Problems. *IEEE Transactions on Circuits and Systems*, 35(10):1243–1256, October 1988.

[25] R. C. Ward. Numerical Computation of the Matrix Exponential with Accuracy Estimate. *SIAM J. Numer. Anal.*, 14(4):600–610, 1977.

[26] *Wolfram Research inc., Mathematica.* Champaign, Illinois, Version 7.0, 2008.