

## Effects of round-to-nearest and stochastic rounding in the numerical solution of the heat equation in low precision<sup>†</sup>

M. Croci<sup>‡</sup> and M. B. Giles<sup>§</sup>

[Received on March 7, 2022]

Motivated by the advent of machine learning, the last few years have seen the return of hardware-supported low-precision computing. Computations with fewer digits are faster and more memory and energy efficient, but can be extremely susceptible to rounding errors. As shown by recent studies into reduced-precision climate simulations, an application that can largely benefit from the advantages of low-precision computing is the numerical solution of partial differential equations (PDEs). However, a careful implementation and rounding error analysis are required to ensure that sensible results can still be obtained.

In this paper we study the accumulation of rounding errors in the solution of the heat equation, a proxy for parabolic PDEs, via Runge-Kutta finite difference methods using round-to-nearest (RtN) and stochastic rounding (SR). We demonstrate how to implement the scheme to reduce rounding errors and we derive *a priori* estimates for local and global rounding errors. Let  $u$  be the unit roundoff. While the worst-case local errors are  $O(u)$  with respect to the discretization parameters (mesh size and timestep), the RtN and SR error behavior is substantially different. In fact, the RtN solution always stagnates for small enough  $\Delta t$ , and until stagnation the global error grows like  $O(u\Delta t^{-1})$ . In contrast, we show that the leading-order errors introduced by SR are zero-mean, independent in space and mean-independent in time, making SR resilient to stagnation and rounding error accumulation. In fact, we prove that for SR the global rounding errors are only  $O(u\Delta t^{-1/4})$  in 1D and are essentially bounded (up to logarithmic factors) in higher dimensions.

**Keywords:** Rounding error analysis, floating-point arithmetic, stochastic rounding, Runge-Kutta methods, low-precision PDE solvers, finite differences.

**AMS subject classifications:** 65G50, 65G30, 65M06, 65M12, 65M15, 65M22, 65Y99, 65C20.

### 1. Introduction

With the increase of machine learning and GPU acceleration applications as a driving force, the last few years have seen the rise of hardware-supported low-precision computing. In fact, support for the IEEE 754 (2008) standard fp16 format and the Google brain floating-point format bfloat16 support is growing among major hardware manufacturers<sup>1</sup>. These 16-bit formats allow for much faster operations with much lower memory, bandwidth, and energy requirements in exchange for a reduction

<sup>†</sup>**Funding:** This research is supported by the ICONIC EPSRC Programme Grant (EP/P020720/1).

<sup>‡</sup>Oden Institute, University of Texas at Austin, Austin, TX, USA, ([matteo.croci@austin.utexas.edu](mailto:matteo.croci@austin.utexas.edu)).

<sup>§</sup>Mathematical Institute, University of Oxford, Oxford, UK, ([mike.giles@maths.ox.ac.uk](mailto:mike.giles@maths.ox.ac.uk)).

<sup>1</sup>The following information is currently available on the manufacturers' respective websites:

The fp16 format is currently supported by AMD (Instinct) and NVIDIA (Pascal, Volta and Ampere) GPUs, ARM (Armv8-A Neon) chips and partially by Intel (Xeon) chips. The bfloat16 format is supported by Intel NNP processors and FPGAs, Google TPUs, and by ARM (Armv8-A) chips. AMD has added support for bfloat16 in its ROCm libraries and Intel has announced that the future Cooper Lake and Sapphire Rapids CPUs will also support the format.

in precision and/or range and a consequential increase in rounding error. In this context, the much larger unit roundoff makes traditional worst-case rounding error bounds too large and thus uninformative, and a naïve implementation can cause numerical algorithms to lose all accuracy.

Nevertheless, careful algorithmic design can still yield extremely fast and relatively accurate results. For this reason, reduced- or mixed-precision algorithms have become very popular in numerical linear algebra (Abdelfattah *et al.*, 2021; Higham & Mary, 2019) and machine learning (Gupta *et al.*, 2015; Na *et al.*, 2017; Ortiz *et al.*, 2018; Wang *et al.*, 2018; Mellempudi *et al.*, 2019) communities. A substantial research effort is also currently spent to reduce the precision in climate and weather model simulations without affecting accuracy. The results are very promising and show large performance gains with negligible accuracy loss (Vána *et al.*, 2017; Düben *et al.*, 2017; Klöwer *et al.*, 2021), and show how numerical algorithms for the solution of PDEs can significantly benefit from a clever reduced-precision implementation.

In this paper we study the accumulation of rounding errors in the solution of the heat equation in  $d \in \{1, 2, 3\}$  dimensions using either round-to-nearest (RtN) or stochastic rounding (SR). SR was first conceptualized by Forsythe (1950, 1959), and by Hull & Swenson (1966) and consists in randomly rounding a number  $x$  in the floating-point range towards minus or plus infinity with different probabilities so that the resulting rounded quantity,  $\text{sr}(x)$ , ( $\text{sr}(x)$  is the SR function, to be defined later in Section 2) is exact in expectation, i.e.  $\mathbb{E}[\text{sr}(x)] = x$ . SR has two disadvantages with respect to RtN: since rounding away from the nearest representable number is possible, it can produce up to twice the error. Furthermore, it requires random number generation, which is likely to be more expensive. Nevertheless, the use of SR comes with its advantages. For an overview of SR hardware and software implementations, analysis, and applications we refer the reader to the review article by Croci *et al.* (2022).

First, the roundoff errors generated by stochastic rounding are random variables, making SR the first rounding mode to be provably amenable to probabilistic rounding error analysis (Connolly *et al.*, 2021). In contrast, the probabilistic treatment of deterministic rounding errors is only a modelling device (Von Neumann & Goldstine, 1947; Henrici, 1962, 1963; Hull & Swenson, 1966; Higham & Mary, 2019) and the smaller probabilistic bounds might not hold for RtN in some scenarios where systematic accumulation of rounding errors occurs (Higham & Mary, 2019).

Second, SR is effective at avoiding stagnation, a phenomenon that occurs e.g. when a (relatively) small, exactly representable number  $\varepsilon$  is added to a large, exactly representable number  $x$ . In RtN mode, if  $\varepsilon$  is smaller than half the gap between  $x$  and its representable neighbours, then the quantity  $x + \varepsilon$  is rounded back to  $x$  and the term  $\varepsilon$  is neglected. With SR this phenomenon does not occur since  $\mathbb{E}[\text{sr}(x + \varepsilon)] = x + \varepsilon$  and the addition is not ignored.

For both of these reasons, SR in low precision has become popular in the machine learning literature since it can yield higher accuracy results at a lower cost when training neural networks (Gupta *et al.*, 2015; Essam *et al.*, 2017; Na *et al.*, 2017; Ortiz *et al.*, 2018; Wang *et al.*, 2018; Mellempudi *et al.*, 2019). The effect of SR in the solution of ordinary and partial differential equations (ODEs and PDEs), and neural ODEs in low precisions has also been investigated experimentally (Hopkins *et al.*, 2020; Fasi & Mikaitis, 2021a; Paxton *et al.*, 2021), showing that SR can still yield accurate results in situations where RtN fails to produce meaningful results. SR is also used in neuromorphic computing (Davies *et al.*, 2018; Sebastian & Mayr, 2018).

While SR resilience to stagnation is now well understood (Connolly *et al.*, 2021), only limited theoretical results are present in the literature about the behavior of SR in common numerical algorithms. The only theoretical work that we are aware of is the recent paper by Connolly *et al.* (2021) in which the authors study the effects of SR in linear algebra operations and show how SR fits within recent research on probabilistic rounding error analysis (Higham & Mary, 2019; Ipsen & Zhou, 2020; Higham

& Mary, 2020). In the context of ODEs, the only applicable results that we are aware of come from the 1960s and the classical work of Henrici (1962, 1963) and the more recent paper by Arató (1983). While SR is not mentioned, Henrici and Arató build on the assumption that at each timestep a local rounding error is introduced, which is treated as an independent, almost surely (a.s.) bounded random variable. A similar assumption is used by Jézéquel (1995), who derives estimates for the first and second moment of the componentwise error caused by deterministic rounding modes in the solution of the heat equation. However, the results in Jézéquel (1995) do not improve on the worst-case linear error growth rate that is typical of deterministic analyses of ODE solvers (Henrici, 1962, 1963). We are not aware of any study in the literature in which sharper probabilistic rounding error estimates are derived in a context specific to PDEs.

It is worth noting that SR hardware support is currently scarce, yet growing due to its popularity in the machine learning and artificial intelligence community. SR is currently used by Intel neuromorphic chip Loihi (Davies *et al.*, 2018) and will be supported in the future neuromorphic supercomputing architecture SpiNNaker2 (Sebastian & Mayr, 2018). The new Graphcore Intelligence Processing Unit (IPU), designed for machine learning and artificial intelligence applications, currently supports SR (Jia *et al.*, 2019). Patents by Intel (Kaul *et al.*, 2018), AMD (Loh, 2019), NVIDIA (Alben *et al.*, 2019), IBM (Bradbury *et al.*, 2017a,b), and other companies (Kanter & Bar, 2010; Henry & Reed, 2018; Lifsches, 2020) have been filed describing possible SR implementations, suggesting that more chips supporting SR might be in the works. Finally, Higham & Pranesh (2019) have developed a low precision emulator which supports stochastic rounding “because of its increasing adoption in machine learning and the need for better understanding of its behavior”. Other low-precision emulators supporting SR are available (cf. Fousse *et al.* (2007); Zhang *et al.* (2019); Fasi & Mikaitis (2021b)), see Section 7(d) in Croci *et al.* (2022) for an overview.

In this paper, we make the following new contributions:

- We analyze how rounding errors affect the numerical solution of the heat equation in 1D, 2D and 3D on the unit box. We consider a standard second-order finite-difference scheme in space and an arbitrary Runge-Kutta method in time, and we show how the numerical scheme should be implemented to reduce the effect of rounding errors and fully exploit the resilience to stagnation of SR. The key element is the use of the delta form of the timestepping scheme. While the delta form is a well-known technique in the numerical analysis and computational fluid dynamics community (cf. Section 12.6 in Lomax *et al.* (2013), and Chapter IV.8 in Hairer & Wanner (1996)), we are unaware of our rigorous analysis of its benefits. Our theory and experiments establish and confirm its efficacy in a probabilistic setting.
- Let  $u$  be the unit roundoff. We show that local rounding errors (introduced at each timestep) are essentially  $O(u)$  with respect to the discretization parameters and are neither independent across timesteps, nor can they be modelled as such. Previous work on probabilistic rounding error analysis for vector ODEs (Henrici, 1962, 1963; Arató, 1983) thus cannot be generalized to the parabolic PDE case since it relies on the local rounding errors to be  $O(u\Delta t^p)$  for some  $p \geq 2$  and independent across timesteps.
- We derive *a priori* error bounds for both local and global rounding errors. The local error estimates are worst-case bounds and hold for both RtN and SR. The global rounding error estimates depend on the dimension and are specific to the rounding mode used. In the latter case strong error bounds for the expected discrete  $L^2$  and  $L^\infty$  norms are provided. The key new result is that while it is known that for RtN the global error grows like  $O(u\Delta t^{-1})$  in  $d$ -dimensions (Henrici, 1962,

1963), for SR it is only  $O(u\Delta t^{-1/4})$  in 1D and essentially  $O(u)$  (up to logarithmic factors) for  $d > 1$ . This is a much stronger result than what is achieved for numerical linear algebra computations, where SR was shown to only reduce the global rounding error from  $O(nu)$  to  $O(\sqrt{nu})$ , ( $n$  here is the problem size) (Connolly *et al.*, 2021). In contrast with numerical linear algebra, when solving the heat equation SR is therefore sufficient to achieve very good accuracy without resorting to other rounding error reducing techniques such as compensated summation; cf. (Higham, 2002, Section 4.3). Numerical results support our findings.

This paper is structured as follows. In Section 2 we give some background information, we describe some basic properties of SR and we introduce the numerical scheme used. In Section 3 we describe how the numerical scheme should be implemented to reduce rounding errors and we derive worst-case local rounding error bounds for both the “naïve” and more accurate implementation. We conclude the section by showing that with RtN the numerical scheme always stagnates for small enough  $\Delta t$ . In Section 4 we combine the local results into global estimates for the global rounding error. We support our findings with numerical results presented in Section 5 and we give some concluding remarks in Section 6.

## 2. Preliminaries

### 2.1 Finite precision arithmetic and types of rounding

In this paper we only consider the base-2 normalized floating-point number systems presented in Table 1. We define a floating-point number system as in the book by Higham (2002):

DEFINITION 2.1 (Section 2.1 in Higham (2002)) A floating-point number system  $F \subset \mathbb{R}$  is a subset of the real numbers whose elements have the form

$$y = \pm m \times \beta^{e-t}, \quad (2.1)$$

where  $\beta, m, t \in \mathbb{N}$  and  $e \in \mathbb{Z}$ . The quantity  $\beta$  is the base ( $\beta = 2$  in this paper),  $t$  is the precision,  $e$  with  $e_{\min} \leq e \leq e_{\max}$  is the exponent and  $m$  with  $0 \leq m \leq \beta^t - 1$  is the significand.

Format	$u$	$x_{\min}$	$x_{\max}$	$t$	exponent bits
bfloat16	$3.91 \times 10^{-3}$	$1.18 \times 10^{-38}$	$3.39 \times 10^{38}$	8	8
fp16	$4.88 \times 10^{-4}$	$6.10 \times 10^{-5}$	$6.55 \times 10^4$	11	5
fp32 (single)	$5.96 \times 10^{-8}$	$1.18 \times 10^{-38}$	$3.40 \times 10^{38}$	24	8
fp64 (double)	$1.11 \times 10^{-16}$	$2.22 \times 10^{-308}$	$1.80 \times 10^{308}$	53	11

Table 1: *Parameters of the floating-point number systems mentioned in the paper. Here  $u = 2^{-t}$  is the unit roundoff,  $x_{\min}$  is the smallest normalized positive number,  $x_{\max}$  is the largest finite number and  $t$  is the precision. Note that bfloat16 and fp32 have the same exponent bits.*

REMARK 2.1 For the sake of simplicity, the theory presented in this paper ignores effects related to the range of  $F$  (such as underflow/overflow). The numerical results of Section 5, however, do account for such behavior.

A rounding function is a function  $\text{fl} : \mathbb{R} \rightarrow F$ . In what follows, we work under the following standard

floating-point error model (cf. Chapter 2 of [Higham \(2002\)](#)):

$$\text{fl}(a \text{ op } b) = (a \text{ op } b)(1 + \delta), \quad |\delta| < 2u, \quad \text{op} \in \{+, -, \times, \backslash\}, \quad (2.2)$$

where  $u$  is the unit roundoff and the factor of 2 is to account for any type of rounding (not just round-to-nearest (RtN)). In this paper, we denote with  $\text{fl}$  the standard RtN function and with  $\text{sr}$  the SR function, defined as follows.

**DEFINITION 2.2** (Stochastic rounding function) Let  $x \in \mathbb{R}$  and  $a, b \in F$  with  $a \leq x \leq b$  to be the numbers of  $F$  that are adjacent to  $x$ . Set  $I = b - a$ . Then the SR function  $\text{sr}(x)$  is the function that rounds  $x$  to  $a$  with probability  $\rho = (b - x)/I$  and to  $b$  with probability  $1 - \rho = (x - a)/I$ .

The following result, proven by [Connolly et al. \(2021\)](#), shows that the roundoff errors arising from SR are mean-independent zero-mean random variable whose absolute value is bounded by  $2u$ .

**LEMMA 2.1** (cf. Lemmas 4.4 and 4.5 in [Connolly et al. \(2021\)](#)) The SR function can be written as  $\text{sr}(x) = x(1 + \delta)$ , where  $|\delta| < 2u$ . Furthermore, the roundoff errors generated by SR are mean-independent zero-mean random variables, i.e. they satisfy

$$\mathbb{E}[\delta_i | \mathcal{F}_{i-1}] = \mathbb{E}[\delta_i] = 0, \quad \forall i, \quad \text{where} \quad \mathcal{F}_i = \{\delta_1, \dots, \delta_i\}, \quad (2.3)$$

assuming that  $\delta_1, \dots, \delta_i$  are generated in this order.

**REMARK 2.2** To simplify the notation, from now on we set  $v = u$  for RtN and  $v = 2u$  for SR so that we can just write  $v$  when presenting results that apply to both. Furthermore, we use the following bounds valid for  $n v < 1$  ([Higham, 2002](#), Lemma 3.1),

$$\prod_{i=1}^m (1 + \delta_i)^{\rho_i} = 1 + \theta_m, \quad \text{with} \quad |\theta_m| \leq \gamma_m \text{ (RtN and SR)}, \quad |\theta_m| \leq \tilde{\gamma}_m \text{ with high probability (SR only)}$$

where  $\delta_i$  are roundoff errors generated with the same rounding mode,  $\rho_i = \pm 1$  and

$$\gamma_m = \frac{m v}{1 - m v}, \quad \tilde{\gamma}_m = \exp\left(\frac{2c\sqrt{m}u + 4mu^2}{1 - 2u}\right) - 1 = 2c\sqrt{m}u + O(u^2).$$

Here  $c > 0$  is an arbitrary constant that can typically be taken to be  $O(1)$  ([Connolly et al., 2021](#)). The  $\gamma_m$  bound is a classical result which always holds for both RtN and SR ([Higham, 2002](#)), while the  $\tilde{\gamma}_m$  bound for SR is a stochastic bound (for SR  $\theta_m$  is a random variable) that was proved recently ([Connolly et al., 2021](#)) and it only holds with high probability. In what follows we only use  $\gamma_m$  for simplicity, but it can be replaced with  $\tilde{\gamma}_m$  in the SR case, yielding smaller constants.

## 2.2 Numerical discretization of the heat equation

In this paper, we investigate the effects of rounding errors in the numerical solution of the heat equation with non-zero forcing term:

$$\begin{cases} \dot{u}(t, \mathbf{x}) - \Delta u(t, \mathbf{x}) = f(t, \mathbf{x}), & \mathbf{x} \in D = (0, 1)^d, \quad t \in (0, T], \\ u(0, \mathbf{x}) = u_0(\mathbf{x}) & \mathbf{x} \in D = (0, 1)^d, \\ u(t, \mathbf{x}) = g(\mathbf{x}) & \mathbf{x} \in \partial D, \quad t \in (0, T], \end{cases} \quad (2.4)$$

where  $0 \neq f \in C^0(\bar{D} \times [0, T])$  and  $u_0 \in C^2(\bar{D})$  with  $u_0|_{\partial D} = g$ . We choose time-independent boundary conditions for simplicity.

Let  $\mathbf{i} \in \mathbb{N}^d$  be a multi-index. We first discretize (2.4) in space at  $K + 1$  discrete equispaced points  $\{\mathbf{x}_i\}$  in each spatial direction with  $\mathbf{x}_i = \mathbf{i}h$  and  $h = 1/K$  using the standard finite difference scheme

$$\dot{U}_i(t) - h^{-2} \sum_{j=1}^d (U_{i+\mathbf{e}_j}(t) - 2U_i(t) + U_{i-\mathbf{e}_j}(t)) = f_i(t), \quad t \in [0, T], \quad (2.5)$$

where  $U_i(t) \approx u(t, \mathbf{x}_i)$ ,  $f_i(t) = f(t, \mathbf{x}_i)$ , and  $\mathbf{e}_j \in \mathbb{R}^d$  for  $j = 1, \dots, d$  are the canonical basis vectors. This can be rewritten in matrix form as:

$$\dot{\mathbf{U}}(t) + \mathbf{A}\mathbf{U}(t) = \mathbf{f}(t), \quad t \in (0, T], \quad (2.6)$$

where  $\mathbf{A} \in \mathbb{R}^{(K-1)^d \times (K-1)^d}$  is the standard finite difference matrix of the  $d$ -dimensional Laplacian and  $\mathbf{f}(t)$  incorporates the boundary conditions.

The eigenvalues and eigenfunctions of the continuous operator  $-\Delta$  are (Courant & Hilbert, 1953, Chapter VI)

$$\lambda_{\mathbf{k}} = \pi^2 \|\mathbf{k}\|_2^2, \quad v_{\mathbf{k}}(\mathbf{x}) = \prod_{j=1}^d \left( \sqrt{2} \sin(\pi \mathbf{k}_j x_j) \right), \quad (2.7)$$

where  $\mathbf{k} \in \mathbb{N}_0^d$  is a multi-index. Their discrete version, the eigenpairs of  $\mathbf{A}$ , are (Kuttler, 1974)

$$\lambda_{\mathbf{k}}^h = 4K^2 \sum_{j=1}^d \sin^2 \left( \frac{\pi \mathbf{k}_j}{2K} \right), \quad v_{\mathbf{k},i}^h = \prod_{j=1}^d \left( \sin \left( \pi i_j \frac{\mathbf{k}_j}{K} \right) \right). \quad (2.8)$$

Here we have indexed the  $\mathbf{k}$ -th eigenvector  $\mathbf{v}_{\mathbf{k}}^h$  of  $\mathbf{A}$  as if it were a  $d$ -dimensional tensor with entries indexed by the multi-index  $\mathbf{i} \in \mathbb{N}_0^d$ . Both  $\mathbf{i}$  and  $\mathbf{k}$  satisfy  $\mathbf{1} \leq \mathbf{i}, \mathbf{k} \leq (K-1)\mathbf{1}$  entrywise, where  $\mathbf{1} \in \mathbb{R}^d$  is the vector of all ones. Independently from the notation adopted, we will always consider  $\mathbf{v}_{\mathbf{k}}$  to be a vector in  $\mathbb{R}^{(K-1)^d}$ . Note that we picked the scaling of the eigenvectors so that

$$\|\mathbf{v}_{\mathbf{k}}^h\|_2^2 = \left( \frac{K}{2} \right)^d, \quad \|\mathbf{v}_{\mathbf{k}}^h\|_{\infty} \leq 1, \quad \forall \mathbf{k}, \quad (2.9)$$

where the norms are to be intended as vector norms. The discrete eigenvectors are orthogonal so that if  $\mathbf{k} \neq \tilde{\mathbf{k}}$  or if  $\mathbf{i} \neq \tilde{\mathbf{i}}$  we have that  $(\mathbf{v}_{\mathbf{k},j}^h, \mathbf{v}_{\tilde{\mathbf{k}},j}^h) = (\mathbf{v}_{\mathbf{k},i}^h, \mathbf{v}_{\tilde{\mathbf{k}},i}^h) = 0$ . Here  $(\cdot, \cdot)$  denotes the vector inner product.

By expanding  $\mathbf{U}(t)$  in terms of the eigenvectors of  $\mathbf{A}$  we can decouple the ODE system (2.6) into many scalar ODEs: set  $\mathbf{U}(t) = \sum_{\mathbf{k}=1}^{k=(K-1)\mathbf{1}} a_{\mathbf{k}}(t) \mathbf{v}_{\mathbf{k}}^h$ , and  $\mathbf{f}(t) = \sum_{\mathbf{k}=1}^{k=(K-1)\mathbf{1}} f_{\mathbf{k}}(t) \mathbf{v}_{\mathbf{k}}^h$ , plug these into (2.6) and multiply by  $\mathbf{v}_{\tilde{\mathbf{k}}}^h$  to obtain

$$\dot{a}_{\tilde{\mathbf{k}}}(t) + \lambda_{\tilde{\mathbf{k}}}^h a_{\tilde{\mathbf{k}}}(t) = f_{\tilde{\mathbf{k}}}(t), \quad t \in (0, T]. \quad (2.10)$$

We now discretize (2.6) in time by applying a generic Runge-Kutta (RK) method with absolute stability function  $S(z)$ ,  $z \in \mathbb{C}$ ; see (Butcher, 2016, Chapter 3). Note that for explicit RK methods  $S(z)$  is always a polynomial while for implicit methods  $S(z)$  is a rational function. In both cases,  $S(z)$  can always be written in the form  $S(z) = 1 + z\tilde{S}(z)$  where  $\tilde{S}(0) = 1$ ; see (Butcher, 2016, Lemma 351A). We

discretize  $\mathbf{U}(t)$  at times  $t^n = n\Delta t$  for  $n = 0, \dots, N = T/\Delta t$  (assuming  $T/\Delta t$  is integer). By applying the RK scheme to (2.10) and by summing over all eigenmodes we obtain the following scheme,

$$\begin{aligned}\mathbf{U}^{n+1} &= S(-\Delta t A) \mathbf{U}^n + \Delta t \mathbf{F}^n = \mathbf{U}^n + (-\Delta t \tilde{S}(-\Delta t A) A \mathbf{U}^n + \Delta t \mathbf{F}^n) \\ &= \mathbf{U}^n + (\Delta \mathbf{U})^n,\end{aligned}\quad (2.11)$$

$$\text{where } \mathbf{F}^n = \sum_{j=1}^s Q_j(-\Delta t A) \mathbf{f}_j^n = \sum_{j=1}^s Q_j(-\Delta t A) \mathbf{f}(t^n + c_j \Delta t). \quad (2.12)$$

Here, all  $Q_j(z)$  are rational functions (or polynomials if the method is explicit) that satisfy  $\sum_{j=1}^s Q_j(z) = \tilde{S}(z)$ . The  $c_j \in \mathbb{R}$  are the nodes of the RK method and  $s \in \mathbb{N}_0$  is the number of stages of the method, e.g. see (Butcher, 2016, Chapter 3). The two equations on the right in (2.11) are the *delta form* of the scheme (Lomax *et al.*, 2013, Section 12.6), (Hairer & Wanner, 1996, Chapter IV.8), in which  $(\Delta \mathbf{U})^n$  is computed first and then added to  $\mathbf{U}^n$ . As we will see in the next section, evaluating the scheme in its delta form helps to reduce rounding errors.

### 3. Local rounding error analysis

When working in finite precision, the numerical scheme (2.11) cannot be evaluated exactly. What we instead compute is

$$\hat{\mathbf{U}}^{n+1} = S(-\Delta t A) \hat{\mathbf{U}}^n + \Delta t \mathbf{F}^n + \boldsymbol{\varepsilon}^n, \quad (3.1)$$

where  $\boldsymbol{\varepsilon}^n \in \mathbb{R}^{(K-1)^d}$  is the local rounding error and it incorporates all rounding errors introduced at timestep  $n$  for a generic implementation. When the numerical scheme is implemented in delta form, we write instead

$$\hat{\mathbf{U}}^{n+1} = \hat{\mathbf{U}}^n + (\Delta \hat{\mathbf{U}})^n + \boldsymbol{\varepsilon}^n + \boldsymbol{\Theta}^n. \quad (3.2)$$

To properly define  $\boldsymbol{\Theta}^n$  and  $\boldsymbol{\varepsilon}^n$ , let  $\tilde{\boldsymbol{\Theta}}^n$  be the local rounding error arising from the computation of  $(\Delta \hat{\mathbf{U}})^n$ , then we have under the standard floating-point error model (2.2),

$$\begin{aligned}\hat{\mathbf{U}}_j^{n+1} - \left( \hat{\mathbf{U}}_j^n + (\Delta \hat{\mathbf{U}})_j^n \right) &= (1 + \delta_j) \left( \hat{\mathbf{U}}_j^n + ((\Delta \hat{\mathbf{U}})_j^n + \tilde{\boldsymbol{\Theta}}_j^n) \right) - \left( \hat{\mathbf{U}}_j^n + (\Delta \hat{\mathbf{U}})_j^n \right) \\ &= \left( \tilde{\boldsymbol{\Theta}}_j^n + \delta_j (\Delta \hat{\mathbf{U}})_j^n + \delta_j \tilde{\boldsymbol{\Theta}}_j^n (\Delta \hat{\mathbf{U}})_j^n \right) + \delta_j \hat{\mathbf{U}}_j^n = (\boldsymbol{\Theta}_j^n) + \boldsymbol{\varepsilon}_j^n,\end{aligned}\quad (3.3)$$

where  $\boldsymbol{\varepsilon}_j^n = \delta_j \hat{\mathbf{U}}_j^n$  essentially captures the leading order error term arising from the addition of  $\hat{\mathbf{U}}^n$  and  $(\Delta \hat{\mathbf{U}})^n$ . The distinction between  $\boldsymbol{\varepsilon}^n$  and  $\boldsymbol{\Theta}^n$  is important: in Sections 4 and 5 we show how the  $\boldsymbol{\varepsilon}^n$  term dominates and is much easier to analyse than  $\boldsymbol{\Theta}^n$ . Note that  $\boldsymbol{\varepsilon}^n$  is typically different from  $\boldsymbol{\varepsilon}^n + \boldsymbol{\Theta}^n$  since the errors come from different implementations. In fact, we show in Section 5 that the delta form local error is typically smaller.

Our final aim is to derive an estimate for the global rounding error, defined as a norm of the quantity

$$\mathbf{E}^n = \hat{\mathbf{U}}^n - \mathbf{U}^n, \quad (3.4)$$

which we will call the global rounding error vector at timestep  $n$ .



REMARK 3.1 For completeness, we define  $\mathcal{E}^{-1}, \boldsymbol{\varepsilon}^{-1}$  to be the local errors due to the rounding of the initial condition. They satisfy  $\mathcal{E}^{-1} \equiv \boldsymbol{\varepsilon}^{-1} := \widehat{\mathbf{U}}^0 - \mathbf{U}^0$ .

In this section we derive worst-case error bounds for the local errors  $\mathcal{E}^n, \boldsymbol{\varepsilon}^n$  and  $\Theta^n$  (valid for both RtN and SR) and we explain why the RtN solution is always affected by stagnation for small enough  $\Delta t$ . We then use the local error results to obtain a bound for the global rounding error in Section 4.

### 3.1 Matrix-vector products in the delta form implementation

When working with the delta form, we assume that the matrix-vector product  $-A\widehat{\mathbf{U}}^n$  is computed via first order differences as follows (using the same notation as in (2.5)).

$$(-A\widehat{\mathbf{U}}^n)_i = h^{-2} \sum_{j=1}^d [(\widehat{\mathbf{U}}_{i+\mathbf{e}_j}^n - \widehat{\mathbf{U}}_i^n) - (\widehat{\mathbf{U}}_i^n - \widehat{\mathbf{U}}_{i-\mathbf{e}_j}^n)]. \quad (3.5)$$

This approach is much less prone to rounding errors than using a direct second order difference (i.e. made of terms  $\widehat{\mathbf{U}}_{i+\mathbf{e}_j}^n - 2\widehat{\mathbf{U}}_i^n + \widehat{\mathbf{U}}_{i-\mathbf{e}_j}^n$ ). The following two theorems explain why. The first theorem is a generalization of Sterbenz lemma (cf. Theorem 2.5 in (Higham, 2002)):

THEOREM 3.1 If  $a, b, c \in \mathbb{R}$  are exactly represented in floating-point arithmetic, and

$$\max(|a-b|, |b-c|) \leq \min(|a|, |b|, |c|) \quad (3.6)$$

then  $a-b$  and  $b-c$  are both computed exactly as long as subnormal numbers are supported or underflow does not occur.

*Proof.* Let  $\bar{e} \in \mathbb{Z}$  be the largest integer so that  $a, b, c$  can be written as  $a = k \times 2^{\bar{e}}, b = l \times 2^{\bar{e}}, c = m \times 2^{\bar{e}}$ , where  $k, l, m \in \mathbb{Z}$ . Hence,

$$a-b = (k-l) \times 2^{\bar{e}}, \quad b-c = (l-m) \times 2^{\bar{e}}. \quad (3.7)$$

The condition on  $a-b$  implies that  $|k-l| \leq \min(|k|, |l|, |m|)$  and therefore (if subnormal numbers are supported or barring underflow)  $(k-l) \times 2^{\bar{e}}$  has an exact floating-point representation, and the same applies to  $(l-m) \times 2^{\bar{e}}$ .  $\square$

Depending on which one of the three statements in the following theorem holds, we have a rounding error arising from the term  $-\Delta t A\widehat{\mathbf{U}}^n$  of order  $O(u\Delta t/h^2)$  (Statement 1),  $O(u\Delta t/h)$  (Statement 2) or  $O(u\Delta t)$  (Statement 3). In the latter two cases, the error is reduced by an  $O(h)$  and  $O(h^2)$  factor respectively, and decreases with  $\Delta t$ , thus explaining why the evaluation of the matrix-vector product via (3.5) is much more accurate.

THEOREM 3.2 Let  $\nu < 1/4$  and consider the evaluation of

$$\sum_{j=1}^d [(\widehat{\mathbf{U}}_{i+\mathbf{e}_j}^n - \widehat{\mathbf{U}}_i^n) - (\widehat{\mathbf{U}}_i^n - \widehat{\mathbf{U}}_{i-\mathbf{e}_j}^n)]. \quad (3.8)$$

The following statements hold:

1. If for all  $i, n$   $|\widehat{\mathbf{U}}_i^n| \leq M$ , then the evaluation of both (3.8) and a direct second-order difference produces a rounding error of at most  $4d\gamma_{d+1}M$ . Furthermore, (3.8) is bounded in absolute value by  $4dM$ .



2. If there exists  $\bar{c} > 0$  such that for all  $i, j, n$

$$\max \left( |\hat{\mathbf{U}}_{i+\mathbf{e}_j}^n - \hat{\mathbf{U}}_i^n|, |\hat{\mathbf{U}}_i^n - \hat{\mathbf{U}}_{i-\mathbf{e}_j}^n| \right) \leq \bar{c}h, \quad (3.9)$$

then (3.8) is bounded in absolute value by  $2\bar{c}dh$  and its evaluation produces a rounding error of at most  $2\bar{c}d\gamma_{d+1}h$ .

3. If for all  $i, j, n$  the assumptions of Theorem 3.1 hold with  $a = \hat{\mathbf{U}}_{i+\mathbf{e}_j}^n$ ,  $b = \hat{\mathbf{U}}_i^n$  and  $c = \hat{\mathbf{U}}_{i-\mathbf{e}_j}^n$ , and there exists  $\bar{c} > 0$  such that

$$|(\hat{\mathbf{U}}_{i+\mathbf{e}_j}^n - \hat{\mathbf{U}}_i^n) - (\hat{\mathbf{U}}_i^n - \hat{\mathbf{U}}_{i-\mathbf{e}_j}^n)| \leq \bar{c}h^2, \quad (3.10)$$

then (3.8) is bounded in absolute value by  $\bar{c}dh^2$  and its evaluation produces a rounding error of at most  $\bar{c}d\gamma_d h^2$ .

*Proof.* The upper bounds on the absolute value of the rounding errors arising from the summation in (3.8) are obtained via the triangle inequality so we only need to prove bounds on the rounding errors of each of the individual terms. Evaluation of each term in the sum in (3.8) in finite precision yields,

$$\left( (\hat{\mathbf{U}}_{i+\mathbf{e}_j}^n - \hat{\mathbf{U}}_i^n)(1 + \delta_1) - (\hat{\mathbf{U}}_i^n - \hat{\mathbf{U}}_{i-\mathbf{e}_j}^n)(1 + \delta_2) \right) (1 + \delta_3), \quad (3.11)$$

where  $\delta_1$  and  $\delta_2$  are the rounding errors in the first order differences and  $\delta_3$  is the roundoff coming from the last subtraction. Note that all entries of  $\hat{\mathbf{U}}^n$  are already rounded by construction, hence exactly representable. Under the assumptions of Statement 1 we have a maximum rounding error of  $4M((1 + \nu)^2 - 1)$  for each term in the sum in (3.8). After accounting for the summation error we obtain a rounding error bound of  $4dM((1 + \nu)^{d+1} - 1) \leq 4d\gamma_{d+1}M$ . The same reasoning yields a bound of  $2\bar{c}d\gamma_{d+1}h$  if (3.9) holds. To prove Statement 3, we note that owing to Theorem 3.1 the first order differences are computed exactly, so  $\delta_1 = \delta_2 = 0$ . This gives an error of  $\bar{c}\nu h^2$  for each term in the summation and a total error bounded above by  $\bar{c}d\gamma_d h^2$ , which is the desired result. If we use a second-order difference, we get instead

$$\left( (\hat{\mathbf{U}}_{i+\mathbf{e}_j}^n - 2\hat{\mathbf{U}}_i^n)(1 + \delta_1) + \hat{\mathbf{U}}_{i-\mathbf{e}_j}^n \right) (1 + \delta_2), \quad (3.12)$$

and the only obtainable upper bound for each term in the summation is the one from Statement 1. Here we have assumed that multiplication by 2 is exact in base-2 floating-point arithmetic, although this assumption matters little for the sake of the argument. The result does not change if a different summation order is used for the evaluation.  $\square$

**REMARK 3.2** In the limit  $h \rightarrow 0$ , the assumptions of Statements 2 and 3 in Theorem 3.2 do not necessarily hold since they require that in the limit  $\hat{\mathbf{U}}_i^n = \hat{\mathbf{U}}_{i \pm \mathbf{e}_j}^n$  for all  $i, j, n$ . However, we believe that it is reasonable to assume at least (3.9) for  $u \ll h$ .

### 3.2 Worst-case bounds for the local rounding errors

In what follows, we prove that the components of  $\mathcal{E}^n$ ,  $\boldsymbol{\varepsilon}^n$ ,  $\boldsymbol{\theta}^n$  are bounded so that there exist positive constants  $\mathcal{E}$ ,  $\varepsilon$ ,  $\theta$  such that  $|\mathcal{E}_j^n| \leq \mathcal{E}$ ,  $|\boldsymbol{\varepsilon}_j^n| \leq \varepsilon$ ,  $|\boldsymbol{\theta}_j^n| \leq \theta$  for all  $j, n$ . Note that  $\mathcal{E}$ ,  $\varepsilon$ , and  $\theta$  typically depend on the rounding mode, on the unit roundoff  $u$ , on the ratio  $\lambda = \Delta t/h^2$ , and on the RK scheme used. We show that when the RK method is explicit  $\mathcal{E}$  and  $\varepsilon$  depend on  $\lambda$ , but not on  $\Delta t$  and  $K$ , while  $\theta$  is small for small  $\Delta t$ . Things behave similarly for implicit methods, but the bounds might also grow with  $K$ , depending on the linear solver used. We make the following three assumptions:

- (A1) There exist constants  $C_f, M_f > 0$  such that for all  $(t, \mathbf{x}) \in [0, T] \times D$ , we have  $|f(t, \mathbf{x})| \leq M_f$  and the evaluation of  $f(t, \mathbf{x})$  in finite precision produces

$$\hat{f}(t, \mathbf{x}) = f(t, \mathbf{x}) + (\Delta f)(t, \mathbf{x}), \text{ with } |(\Delta f)(t, \mathbf{x})| \leq C_f v |f(t, \mathbf{x})| \leq C_f M_f v.$$

- (A2) There exists  $M > 0$  such that the quantities  $\hat{\mathbf{U}}^n$  satisfy  $\|\hat{\mathbf{U}}^n\|_\infty \leq M$  for all  $n$ . Note that  $M$  might depend on  $h$  and  $\Delta t$ .

- (A3) We assume that for  $p \geq 0$ ,  $\|\Delta t A \hat{\mathbf{U}}^n\|_\infty \leq c_a(d, \lambda) \Delta t^p$  for some constant  $c_a(d, \lambda) > 0$ , and that the evaluation of  $-\Delta t A \hat{\mathbf{U}}^n$  yields an error bounded above by  $c_a(d, \lambda) \Delta t^p v$ . For instance, this is the case if the hypotheses of Theorem 3.2 are satisfied with  $p = 0$ ,  $p = 1/2$  or  $p = 1$  according to whether Statement 1, 2 or 3 holds.

First, we need an auxiliary result:

**THEOREM 3.3** Let  $\mathbf{v}$  be an arbitrary vector in  $\mathbb{R}^{(K-1)^d}$ . Let  $P(x) = N(x)/W(x)$  be either  $S(x)$ ,  $\tilde{S}(x)$  or  $Q_j(x)$  for any  $j = 1, \dots, s$ . Assume that  $P(x)$  has no poles on the negative real axis. Let  $\mathbf{z}$  solve  $W(-\Delta t A)\mathbf{z} = N(-\Delta t A)\mathbf{v}$  and let  $\hat{\mathbf{z}}$  be the result of computing  $\mathbf{z}$  in finite precision by solving the linear system using a stable solver. Then, there exist constants  $C_0, C > 0$  such that

$$\|\mathbf{z}\|_\infty \leq C_0(s, d, \lambda, K) \|\mathbf{v}\|_\infty, \quad \|\hat{\mathbf{z}} - \mathbf{z}\|_\infty \leq C(s, d, \lambda, K) v \|\mathbf{v}\|_\infty. \quad (3.13)$$

The second bound only holds for  $\mathbf{v}$  sufficiently small. If the RK method is explicit the constants  $C_0$  and  $C$  still depend on  $\lambda$ , but not directly on  $K$  (recall that  $K = 1/h$ ). The bound still holds, albeit with a slightly larger constant, if  $\mathbf{v}$  is also computed inexactly yielding  $\hat{\mathbf{v}}$  with  $\|\hat{\mathbf{v}} - \mathbf{v}\|_\infty \leq c v \|\mathbf{v}\|_\infty$  for some  $c > 0$ .

*Proof.* This is a combination of classical results from Higham (2002), cf. Appendix A.  $\square$

We recall that requiring that a linear system, e.g.  $W\mathbf{z} = \mathbf{y}$ , is solved by using a stable solver means that the computed solution  $\hat{\mathbf{z}}$  satisfies for some  $c_1, c_2 > 0$ , (cf. Higham (2002))

$$(W + \Delta W)\hat{\mathbf{z}} = \mathbf{y} + \Delta \mathbf{y}, \quad \text{where} \quad \|\Delta W\|_\infty \leq c_1 v, \quad \|\Delta \mathbf{y}\|_\infty \leq c_2 v. \quad (3.14)$$

Here we focus on rounding errors, but (3.14) and Theorem 3.3 can easily be extended to account for other sources of error arising, e.g. when an iterative method is used. We can now prove the following theorem:

**THEOREM 3.4** Let Assumptions (A1), (A2) and (A3) and the assumptions of Theorem 3.3 hold. Then, there exist constants  $\bar{C}_1, \bar{C}_2, \bar{C}_3 > 0$  depending on the rounding function used such that, for all  $n$ ,

$$\|\mathcal{E}^n\|_\infty \leq (\bar{C}_1(s, d, \lambda, K)M + \Delta t \bar{C}_2(s, d, \lambda, K)M_f) v \equiv \mathcal{E}, \quad (3.15)$$

$$\|\Theta^n\|_\infty \leq \bar{C}_3(p, s, d, \lambda, K) \Delta t^p v \equiv \theta, \quad \|\mathcal{E}^n\|_\infty \leq M v \equiv \mathcal{E}. \quad (3.16)$$

where  $\lambda = \Delta t/h^2$  and the first two bounds only hold for  $\mathbf{v}$  sufficiently small. If the RK method is explicit  $\bar{C}_1, \bar{C}_2$ , and  $\bar{C}_3$  still depend on  $\lambda$ , but not directly on  $K$  (recall that  $K = 1/h$ ).

*Proof.* The bound on  $\|\mathcal{E}^n\|_\infty$  is a straightforward application of Assumption (A2). We first derive a bound for  $\|\mathcal{E}^n\|_\infty$  and then we bound  $\|\Theta^n\|_\infty$ . Let  $\tilde{\mathbf{U}}^{n+1}$  exactly satisfy (3.1) with  $\mathcal{E}^n$  set to zero:

$$\tilde{\mathbf{U}}^{n+1} = S(-\Delta t A)\hat{\mathbf{U}}^n + \Delta t \mathbf{F}^n, \quad (3.17)$$

i.e.  $\tilde{\mathbf{U}}^{n+1}$  is the result of performing one timestep in exact arithmetic starting from  $\hat{\mathbf{U}}^n$ . We can then estimate  $\mathcal{E}$  by computing an upper bound for the quantity

$$\|\mathcal{E}^n\|_\infty = \|\hat{\mathbf{U}}^{n+1} - \tilde{\mathbf{U}}^{n+1}\|_\infty. \quad (3.18)$$

Let  $\hat{\mathbf{f}}_j^n = \mathbf{f}_j^n + (\Delta \mathbf{f}_j)^n$  for  $j = 1, \dots, s$  be the computed forcing terms due to finite precision and note that by Assumption (A1) we have  $\|(\Delta \mathbf{f}_j)^n\|_\infty \leq C_f M_f v$ . Let  $\mathbf{z}_0 = S(-\Delta t A) \hat{\mathbf{U}}^n$ ,  $\mathbf{y} = \tilde{S}(-\Delta t A)(-\Delta t A \hat{\mathbf{U}}^n)$  and  $\mathbf{z}_j = \Delta t Q_j(-\Delta t A) \hat{\mathbf{f}}_j^n$ , and let  $\hat{\mathbf{z}}_j$  for  $j = 0, \dots, s$  and  $\hat{\mathbf{y}}$  be the result of computing the  $\mathbf{z}_j$  and  $\mathbf{y}$  in finite precision. Owing to Theorem 3.3 and Assumptions (A1)-(A3), we obtain the upper bounds

$$\begin{aligned} \|\hat{\mathbf{z}}_0 - \mathbf{z}_0\|_\infty &\leq C_1 v M, & \|\hat{\mathbf{z}}_j - \mathbf{z}_j\|_\infty &\leq C_2 \Delta t v M_f, & \|\hat{\mathbf{y}} - \mathbf{y}\|_\infty &\leq C_3 \Delta t^p v, \\ \|\mathbf{z}_0\|_\infty &\leq c_1 M, & \|\mathbf{z}_j\|_\infty &\leq c_2 \Delta t M_f, & \|\mathbf{y}\|_\infty &\leq c_3 \Delta t^p, \end{aligned} \quad (3.19)$$

where  $c_1, c_2, c_3, C_1, C_2$  and  $C_3$  are positive constants that depend on  $s, d$  and  $\lambda$  and might depend on  $K$  for implicit methods. We first collect (3.19) into a bound for (3.18). We can write  $\hat{\mathbf{U}}^{n+1}$  as

$$\hat{\mathbf{U}}^{n+1} = \sum_{j=0}^s \hat{\mathbf{z}}_j \circ (1 + \theta_j), \quad \text{where } |\theta_j| \leq \gamma_s \mathbf{1}. \quad (3.20)$$

Here  $\circ$  denotes the entrywise product. The vectors  $\theta_j$  represent the error introduced by summing the  $\hat{\mathbf{z}}_j$  together in finite precision, cf. (Higham, 2002, Chapter 3). Owing to (3.19) and Theorem 3.3, we have

$$\begin{aligned} \|\mathcal{E}^n\|_\infty &= \|\hat{\mathbf{U}}^{n+1} - \tilde{\mathbf{U}}^{n+1}\|_\infty = \left\| \sum_{j=0}^s \hat{\mathbf{z}}_j \circ (1 + \theta_j) - \sum_{j=0}^s \mathbf{z}_j \right\|_\infty \\ &\leq (1 + \gamma_s) \sum_{j=0}^s \|\hat{\mathbf{z}}_j - \mathbf{z}_j\|_\infty + \gamma_s \sum_{j=0}^s \|\mathbf{z}_j\|_\infty \\ &\leq \bar{C}_1(s, d, \lambda, K) M v + \bar{C}_2(s, d, \lambda, K) \Delta t M_f v = \mathcal{E}. \end{aligned} \quad (3.21)$$

This is the desired result. We omit the derivation of a close-form expression for  $\bar{C}_1$  and  $\bar{C}_2$ . The reasoning for  $\|\Theta^n\|_\infty$  is similar. First we use (3.19) to bound

$$\|\Delta \hat{\mathbf{U}}\|_\infty \leq \|\mathbf{y}\|_\infty + \Delta t \sum_{j=1}^s \|\mathbf{z}_j\|_\infty \leq c_3 \Delta t^p + c_2 \Delta t M_f. \quad (3.22)$$

Following the same procedure as the one used to obtain (3.21), we get for  $\|\tilde{\Theta}^n\|_\infty$  (cf. (3.3)),

$$\begin{aligned} \|\tilde{\Theta}^n\|_\infty &\leq (1 + \gamma_s) \|\hat{\mathbf{y}} - \mathbf{y}\|_\infty + (1 + \gamma_s) \sum_{j=1}^s \|\hat{\mathbf{z}}_j - \mathbf{z}_j\|_\infty + \gamma_s \left( \|\mathbf{y}\|_\infty + \sum_{j=1}^s \|\mathbf{z}_j\|_\infty \right) \\ &\leq \tilde{C}_3(s, d, \lambda, K) \Delta t^p v + \tilde{C}_2(s, d, \lambda, K) \Delta t M_f v, \end{aligned} \quad (3.23)$$

for some  $\tilde{C}_3 > 0$  that does not depend on  $K$  for explicit methods. The definition of  $\Theta^n$ , given entrywise in (3.3), implies that

$$\|\Theta^n\|_\infty \leq \|\tilde{\Theta}^n\|_\infty + v \|\Delta \hat{\mathbf{U}}\|_\infty + v \|\tilde{\Theta}^n\|_\infty \|\Delta \hat{\mathbf{U}}\|_\infty. \quad (3.24)$$

Pulling together equations (3.22), (3.23), and (3.24) yields an upper bound for  $\|\Theta^n\|_\infty$  and the desired result. We omit the derivation of a close-form expression for  $\tilde{C}_3$ .  $\square$

REMARK 3.3 The constants  $\bar{C}_1$ ,  $\bar{C}_2$  and  $\bar{C}_3$  in Theorem 3.4 typically grow at least linearly with the condition number of  $W(-\Delta t A)$ , which is typically  $O(\lambda)$ . While it is well-known that  $\lambda$  is bounded above for explicit methods due to stability reasons, this suggests that  $\lambda$  must be kept small for implicit methods as well to avoid losing all accuracy.

### 3.3 Round-to-nearest causes stagnation

As discussed in (Gupta *et al.*, 2015; Mikaitis, 2021; Fasi & Mikaitis, 2021a; Connolly *et al.*, 2021) RtN computations are prone to stagnation. Explaining why this phenomenon occurs with RtN in an ODE/PDE context for small enough  $\Delta t$  is simple. Let  $x$  and  $\varepsilon$  be two exactly representable numbers. Then whenever the magnitude of  $\varepsilon$  is smaller than half the gap between  $x$  and its representable neighbours,  $x + \varepsilon$  is rounded back to  $x$ . In formulas (cf. Lemma 2.1 in Higham (2002)),

$$\text{fl}(x + \varepsilon) = x, \quad \text{whenever} \quad \frac{u}{2}|x| \geq |\varepsilon|. \quad (3.25)$$

We now consider stagnation in our numerical scheme. For simplicity, we only consider a delta form implementation of the RK method. Let  $u(t, \mathbf{x})$  be the solution of (2.4), then stagnation will approximately occur when

$$\frac{u}{2}|u(t_n, \mathbf{x}_i)| \approx \frac{u}{2}|\hat{\mathbf{U}}_i^n| \geq |\Delta \hat{\mathbf{U}}_i^n| = |\hat{\mathbf{U}}_i^{n+1} - \hat{\mathbf{U}}_i^n| \approx \Delta t |\dot{u}(t_n, \mathbf{x}_i)|. \quad (3.26)$$

This shows that  $\hat{\mathbf{U}}_i^n$  will not be updated whenever

$$|u(t_n, \mathbf{x}_i)| \gtrsim 2(\Delta t/u)|\dot{u}(t_n, \mathbf{x}_i)|. \quad (3.27)$$

This gives a condition which is clearly discretization-dependent and shows that stagnation will lead to different discrete solutions for different values of  $\Delta t$  (hence  $h$ ) and the unit roundoff  $u$ . According to (3.27), when a non-zero steady-state solution exists stagnation always occurs, since  $|\dot{u}(t, \mathbf{x})| \rightarrow 0$  as  $t \rightarrow \infty$ . Furthermore, if  $u(0, \mathbf{x}), \dot{u}(0, \mathbf{x}) \neq 0$  the discrete solution will always stagnate at the initial condition for  $\Delta t$  small enough.

Overall, this simple argument already shows that the steady-state solution obtained with RtN is discretization, initial condition and precision dependent. This argument can be made rigorous, but we believe that the fact that stagnation occurs in this setting is both obvious and well known in the community, so we omit further details. We remark that not using the delta form only worsened the stagnation behavior in numerical experimentations.

## 4. Global rounding error analysis

We now derive expressions for the global rounding errors in terms of  $\mathcal{E}$ ,  $\theta$ , and  $\varepsilon$ . From Theorem 3.4 we know that the worst-case error bounds  $|\mathcal{E}_i^n| \leq \mathcal{E}$ ,  $|\boldsymbol{\varepsilon}_i^n| \leq \varepsilon$ , and  $|\Theta_i^n| \leq \theta$  hold. However, more precise bounds can be obtained for SR owing to the following lemma:

LEMMA 4.1 With SR we have that  $\boldsymbol{\varepsilon}_i^n$  is independent from  $\boldsymbol{\varepsilon}_j^n$  for all  $n, i \neq j$ . Furthermore,  $\boldsymbol{\varepsilon}_i^n$  satisfies

$$\mathbb{E}[\boldsymbol{\varepsilon}_i^n] = \mathbb{E} \left[ \boldsymbol{\varepsilon}_i^n \left| \left\{ \{\boldsymbol{\varepsilon}_j^m\}_{j=1}^{j=(K-1)1} \right\}_{m=1}^{n-1} \right. \right] = 0 \text{ for all } n, i.$$

*Proof.* The thesis is a consequence of the definition of  $\boldsymbol{\varepsilon}^n$  and of the statistical properties of the SR roundoff errors, cf. Lemma 2.1.  $\square$

We can now distinguish two different scenarios:

- 1) **Worst-case (RtN):** we can only assume that  $|\boldsymbol{\epsilon}_i^n| \leq \varepsilon$  for all  $n, i$ .
- 2) **Stochastic rounding:** Lemma 4.1 holds.

REMARK 4.1 From now on we just investigate the global error deriving from  $\boldsymbol{\epsilon}^n$  for simplicity. However, the error bounds obtained also apply to the global errors deriving from the other error terms after making the following substitutions:

- Scheme not in delta form (both RtN and SR): simply replace  $\varepsilon$  with  $\boldsymbol{\mathcal{E}}$  in the results from Scenario 1).
- Scheme in delta form, RtN: the total global error is obtained by replacing  $\varepsilon$  with  $\varepsilon + \theta$  in the results from Scenario 1).
- Scheme in delta form, SR: the total global error consists of two additive terms. The former is given by Scenario 2) and the latter is given by replacing  $\varepsilon$  with  $\theta$  in Scenario 1). While our theory does not show that the  $\theta$  error term is negligible for small  $\Delta t$ , the effects of this second error term do not actually appear in numerical experimentation (cf. Section 5).

REMARK 4.2 In this section we assume for simplicity that the initial condition can be represented exactly so that  $\hat{\boldsymbol{U}}^0 = \boldsymbol{U}^0$  and  $\boldsymbol{\mathcal{E}}^{-1} = \boldsymbol{\epsilon}^{-1} = \mathbf{0}$ .

We present a summary of the global error asymptotic rates derived in this section for fixed  $\lambda$  in Table 2. We use both the infinity norm and the discrete  $L^2$  norm, defined as  $\|\cdot\|_{L^2} = K^{-d/2} \|\cdot\|_2$ . The Scenario 1) infinity norm results for  $d > 1$  depend on the assumption that the matrix  $S(-\Delta t A)$  only has non-negative entries. This requirement also implies that a discrete maximum principle holds and is automatically satisfied for Forward Euler and the popular RK4 method in their stability region, for Backward Euler for all  $\lambda$ , and for the Crank-Nicolson method for  $\lambda \leq 1/d$ .

Scenario	Norm	1D	2D	3D
1) RtN	$\infty$	$O(\varepsilon \Delta t^{-1})$	$O(\varepsilon \Delta t^{-1} \ell(\Delta t)^s)$	$O(\varepsilon \Delta t^{-1-s/2})$
1) RtN	$L^2$	$O(\varepsilon \Delta t^{-1})$	$O(\varepsilon \Delta t^{-1})$	$O(\varepsilon \Delta t^{-1})$
2) SR	$\infty$	$O(\varepsilon \Delta t^{-1/4} \ell(\Delta t)^{1/2})$	$O(\varepsilon \ell(\Delta t))$	$O(\varepsilon \ell(\Delta t)^{1/2})$
2) SR	$L^2$	$O(\varepsilon \Delta t^{-1/4})$	$O(\varepsilon \ell(\Delta t)^{1/2})$	$O(\varepsilon)$

Table 2: Asymptotic global rounding error blow-up rates for fixed  $\lambda$  in the infinity and discrete  $L^2$  norm. Here  $\ell(\Delta t) = |\log(\lambda^{-1} \Delta t)|$ . The error measures used for the worst-case scenario are either  $\|\boldsymbol{E}^n\|_\infty$  or  $\|\boldsymbol{E}^n\|_{L^2}$ . The value of  $s$  is 0 if  $S(-\Delta t A)$  only has non-negative entries and  $s = 1$  otherwise. For the SR case the error measures are either  $\mathbb{E}[\|\boldsymbol{E}^n\|_\infty^2]^{1/2}$  or  $\mathbb{E}[\|\boldsymbol{E}^n\|_{L^2}^2]^{1/2}$ .

REMARK 4.3 For ease of notation, in this section we write e.g.  $\sum_{\mathbf{k}}$  in place of  $\sum_{\mathbf{k}=1}^{k=(K-1)1}$  to indicate a sum over all indices of the spatial degrees of freedom.

Some preliminary analysis is needed before we can go over each scenario. We start by subtracting the exact delta form scheme (2.11) from the approximate (3.2) and we set  $\boldsymbol{\Theta}^n = \mathbf{0}$  to get a recursive relation for the global rounding error  $\boldsymbol{E}^n = \hat{\boldsymbol{U}}^n - \boldsymbol{U}^n$  deriving from  $\boldsymbol{\epsilon}^n$ :

$$\boldsymbol{E}^{n+1} = S(-\Delta t A) \boldsymbol{E}^n + \boldsymbol{\epsilon}^n. \quad (4.1)$$

The first step is to expand  $\mathbf{E}^n$  in terms of the eigenvectors of  $A$  (cf. (2.8)). We write  $\mathbf{E}^n = \sum_{\mathbf{k}} a_{\mathbf{k}}^n \mathbf{v}_{\mathbf{k}}^h$  and we plug this into the above to obtain

$$\sum_{\mathbf{k}} a_{\mathbf{k}}^{n+1} \mathbf{v}_{\mathbf{k}}^h = \sum_{\mathbf{k}} a_{\mathbf{k}}^n S(s_{\mathbf{k}}) \mathbf{v}_{\mathbf{k}}^h + \boldsymbol{\varepsilon}^n, \quad (4.2)$$

where  $s_{\mathbf{k}} = -\Delta t \lambda_{\mathbf{k}}^h$ . Multiplying both sides by  $(\mathbf{v}_{\bar{\mathbf{k}}}^h)^T$  we get, due to orthogonality and (2.9),

$$a_{\bar{\mathbf{k}}}^{n+1} = a_{\bar{\mathbf{k}}}^n S(s_{\bar{\mathbf{k}}}) + \left(\frac{2}{K}\right)^d (\boldsymbol{\varepsilon}^n, \mathbf{v}_{\bar{\mathbf{k}}}^h), \quad (4.3)$$

where we write  $(\mathbf{a}, \mathbf{b}) = \mathbf{a}^T \mathbf{b}$  for any two vectors  $\mathbf{a}, \mathbf{b}$ . We now drop the bar in  $\bar{\mathbf{k}}$ . Since  $\mathbf{E}^0 = \mathbf{0}$  by assumption (cf. Remark 4.2), we have  $a_{\mathbf{k}}^0 = 0$  for all  $\mathbf{k}$  and we obtain

$$a_{\mathbf{k}}^n = \left(\frac{2}{K}\right)^d \sum_{m=0}^{n-1} S(s_{\mathbf{k}})^{n-1-m} (\boldsymbol{\varepsilon}^m, \mathbf{v}_{\mathbf{k}}^h). \quad (4.4)$$

Before working through the various correlation assumptions it is useful to prove the following theorem, as it plays an important role in the whole section.

**THEOREM 4.1** Let  $\mathcal{A} = \{z \leq 0 : |S(z)| < 1\}$ ,  $\Lambda = (-4d\lambda, 0)$ , and assume that  $\lambda$  is chosen such that  $(\Lambda \cup \{-4d\lambda\}) \subset \mathcal{A}$ . Further assume that  $S(z)$  has no poles on the negative real axis. Then, there exist constants  $C(d, \lambda) > 0, \bar{C}(d, \lambda), \tilde{C}(d, \lambda) > 0$  such that

$$\sum_{\mathbf{k}} \frac{1}{1 - S(s_{\mathbf{k}})^2} \leq \sum_{\mathbf{k}} \frac{1}{1 - |S(s_{\mathbf{k}})|} \leq \phi_d(\Delta t) = \begin{cases} C(1, \lambda) \Delta t^{-1}, & d = 1, \\ C(2, \lambda) \Delta t^{-1} |\log(\lambda^{-1} \Delta t)|, & d = 2, \\ C(3, \lambda) \Delta t^{-3/2}, & d = 3, \end{cases} \quad (4.5)$$

$$\sum_{\mathbf{k}} \frac{1}{(1 - |S(s_{\mathbf{k}})|)^2} \leq \bar{C}(d, \lambda)^2 \Delta t^{-2}, \quad (4.6)$$

$$K^{-d} \sum_{\mathbf{k}} \frac{(\mathbf{1}, \mathbf{v}_{\mathbf{k}}^h)}{1 - |S(s_{\mathbf{k}})|} \leq \tilde{C}(d, \lambda) \Delta t^{-1}. \quad (4.7)$$

*Proof.* See [Appendix B](#). □

We now derive estimates for the global error for both scenarios.

#### 4.1 Worst-case error (round-to-nearest)

In this case, the only thing we know is that  $|\boldsymbol{\varepsilon}_i^n| \leq \varepsilon$  for all  $n, i$ . We have

$$|a_{\mathbf{k}}^n| \leq 2^d \varepsilon \left(\frac{K-1}{K}\right)^d \sum_{m=0}^{n-1} |S(s_{\mathbf{k}})|^{n-1-m} \leq 2^d \frac{\varepsilon}{1 - |S(s_{\mathbf{k}})|},$$

since for stability  $|S(s_k)| < 1$ . Let  $\mathbf{a}^n \in \mathbb{R}^{(K-1)^d}$  be such that  $(\mathbf{a}^n)_k = a_k^n$ . We then have that (cf. (2.9) and Theorem 4.1)

$$\begin{aligned} \|\mathbf{E}^n\|_\infty &\leq \sum_k |a_k^n| \leq 2^d \varepsilon \sum_k \frac{1}{1 - |S(s_k)|} \leq 2^d \varepsilon \phi_d(\Delta t), \\ \|\mathbf{E}^n\|_{L^2}^2 &= K^{-d} \left(\frac{K}{2}\right)^d \|\mathbf{a}^n\|_2^2 \leq \varepsilon^2 2^d \sum_k \frac{1}{(1 - |S(s_k)|)^2} \leq 2^d \bar{C}(d, \lambda)^2 \varepsilon^2 \Delta t^{-2}, \\ \Rightarrow \|\mathbf{E}^n\|_{L^2} &\leq 2^{d/2} \bar{C}(d, \lambda) \varepsilon \Delta t^{-1}. \end{aligned}$$

Here we used the equality  $\|\mathbf{E}^n\|_2^2 = (K/2)^d \|\mathbf{a}^n\|_2^2$  which comes from the fact that the  $\mathbf{v}_k^h$  are orthogonal. The function  $\phi_d(\Delta t)$  was defined in (4.5):  $\phi_1(\Delta t) = C(1, \lambda) \Delta t^{-1}$ ,  $\phi_2(\Delta t) = C(2, \lambda) \Delta t^{-1} |\log(\lambda^{-1} \Delta t)|$ ,  $\phi_3(\Delta t) = C(3, \lambda) \Delta t^{-3/2}$ .

The infinity norm estimate in the 2D and 3D case can be refined when the matrix  $S(-\Delta t A)$  has only non-negative entries. In this case, a discrete maximum (minimum) principle holds and for an arbitrary sequence of non-negative (respectively non-positive) vectors  $\mathbf{f}^0, \dots, \mathbf{f}^n$  we have that the vector sequence  $\mathbf{W}^n$  satisfying

$$\mathbf{W}^{n+1} = S(-\Delta t A) \mathbf{W}^n + \mathbf{f}^n = \sum_{m=0}^n S(-\Delta t A)^m \mathbf{f}^m, \quad (4.8)$$

with  $\mathbf{W}^0 = \mathbf{0}$  is also non-negative (respectively non-positive). Now, let  $\mathbf{M}^n$  be the solution of (4.1) in which  $\boldsymbol{\varepsilon}^n$  is replaced by  $\varepsilon \mathbf{1}$ . We then have that  $\mathbf{E}^n - \mathbf{M}^n$  and  $\mathbf{M}^n + \mathbf{E}^n$  satisfy equation (4.8) with right hand sides that are non-positive and non-negative respectively, giving  $\mathbf{E}^n - \mathbf{M}^n \leq 0$  and  $\mathbf{M}^n + \mathbf{E}^n \geq 0$ . This gives  $|\mathbf{E}^n| \leq \mathbf{M}^n$  and

$$\begin{aligned} |\mathbf{E}_j^n| &\leq \mathbf{M}_j^n = \varepsilon \left(\frac{2}{K}\right)^d \sum_k \sum_{m=0}^{n-1} S(s_k)^{n-1-m} (\mathbf{1}, \mathbf{v}_k^h) \mathbf{v}_{k,j}^h \leq \varepsilon 2^d K^{-d} \sum_k \frac{(\mathbf{1}, \mathbf{v}_k^h)}{1 - |S(s_k)|} \\ &\leq \varepsilon 2^d \bar{C}(d, \lambda) \Delta t^{-1}, \end{aligned}$$

where we used (2.9) in the second step and Theorem 4.1 in the third. Since the bound is independent from  $j$  this is also a valid bound for  $\|\mathbf{E}^n\|_\infty$ .

We remark that the fact that RtN rounding errors grow like  $O(u \Delta t^{-1})$  is well-known in an ODE context (Henrici, 1962, 1963). This error growth rate was also more recently observed by Jézéquel (1995), where the author studies the propagation of rounding errors in the solution of the heat equation via the forward Euler method with RtN and towards-zero rounding.

#### 4.2 Stochastic rounding error

Here  $\mathbb{E}[\boldsymbol{\varepsilon}_i^n] = 0$  for all  $n, i$ , and  $\mathbb{E}[\boldsymbol{\varepsilon}_i^n \boldsymbol{\varepsilon}_j^k] = 0$  unless  $n = k$  and  $i = j$  due to mean-independence, cf. Lemma 4.1. We therefore have that owing to (2.9),

$$\mathbb{E}[(\boldsymbol{\varepsilon}^m, \mathbf{v}_k^h)^2] = \sum_{i,j} \mathbb{E}[\boldsymbol{\varepsilon}_i^m \boldsymbol{\varepsilon}_j^m] (\mathbf{v}_k^h)_i (\mathbf{v}_k^h)_j = \varepsilon^2 \|\mathbf{v}_k\|_2^2 = \varepsilon^2 \left(\frac{K}{2}\right)^d.$$



This yields,

$$\begin{aligned}\mathbb{E}[\|\mathbf{E}^n\|_2^2] &= \left(\frac{K}{2}\right)^d \mathbb{E}[\|\mathbf{a}^n\|_2^2] = \left(\frac{2}{K}\right)^d \sum_{\mathbf{k}} \sum_m^{n-1} S(s_{\mathbf{k}})^{2(n-1-m)} \mathbb{E}[(\boldsymbol{\varepsilon}^m, \mathbf{v}_{\mathbf{k}}^h)^2] \\ &\leq \varepsilon^2 \sum_{\mathbf{k}} \frac{1}{1 - S(s_{\mathbf{k}})^2} \leq \varepsilon^2 \phi_d(\Delta t), \\ \Rightarrow \mathbb{E}[\|\mathbf{E}^n\|_{L^2}^2]^{1/2} &\leq \varepsilon K^{-d/2} \sqrt{\phi_d(\Delta t)} = \varepsilon \lambda^{-d/4} \Delta t^{d/4} \sqrt{\phi_d(\Delta t)}.\end{aligned}$$

We have again used Theorem 4.1. The bound is  $O(\Delta t^{-1/4})$  in 1D,  $O(|\log(\lambda^{-1} \Delta t)|^{1/2})$  in 2D and  $O(1)$  in 3D for fixed  $\lambda$ .

Obtaining a bound for  $\mathbb{E}[\|\mathbf{E}^n\|_\infty^2]^{1/2}$  is more difficult and requires the combined use of the concept of a martingale, Azuma-Hoeffding inequality and a maximal inequality for sub-Gaussian random variables. We start by defining martingales:

**DEFINITION 4.2 (Martingale)** A sequence of random variables  $\{X_n\}_{n=0}^N$  is a martingale if  $\mathbb{E}[X_n] < \infty$  and  $\mathbb{E}[X_n | X_0, \dots, X_{n-1}] = X_{n-1}$ .

Note that this is not a complete definition of a martingale, but it is sufficient for our purposes; see Section 2.3 in [Kloeden & Platen \(2013\)](#) for a more complete definition. Azuma-Hoeffding inequality reads:

**THEOREM 4.3 (Azuma-Hoeffding inequality ([Hoeffding \(1963\)](#); [Azuma \(1967\)](#)))** For a martingale  $\{X_n\}_{n=0}^N$  with  $|X_n - X_{n-1}| \leq c_n$  a.s.,

$$\mathbb{P}(|X_N - X_0| > t) < 2 \exp\left(-\frac{t^2}{2 \sum_{n=1}^N c_n^2}\right). \quad (4.9)$$

A sub-Gaussian random variable is a random variable for which the tails of its distribution decay as fast as a Gaussian distribution. More formally,

**DEFINITION 4.4 (Sub-Gaussian random variable)** A zero-mean random variable  $X(\omega)$  is sub-Gaussian with proxy variance  $\sigma^2 > 0$  if

$$\mathbb{P}(|X| > t) \leq 2 \exp\left(-\frac{t^2}{2\sigma^2}\right). \quad (4.10)$$

Clearly, the quantity  $X_N - X_0$  in Azuma-Hoeffding inequality is sub-Gaussian. Here we will be interested in controlling the maximum of many sub-Gaussian random variables. The following maximal inequality will be useful:

**THEOREM 4.5** Let  $\{X_i\}_{i=1}^N$  be a collection of sub-Gaussian random variables with variance proxy  $\sigma^2$ . Then

$$\mathbb{E}[\max_i |X_i|^2] \leq 2\sigma^2 \log(2eN). \quad (4.11)$$

The random variables in the collection need not be independent.

*Proof.* Set  $Z = \max_i |X_i|^2$  and let  $P(z)$  and  $P^{-1}(z)$  be the cumulative distribution function and the quantile function of  $Z$  respectively. Then,

$$\mathbb{E}[Z] = \int_0^1 P^{-1}(s) ds = \int_0^1 P^{-1}(1-s) ds. \quad (4.12)$$

Furthermore, since  $\{X_i\}_{i=1}^N$  are sub-Gaussian, we have

$$1 - P(t) = \mathbb{P}(Z > t) \leq \sum_{i=1}^N \mathbb{P}(|X_i| > t) \leq 2N \exp\left(-\frac{t}{2\sigma^2}\right). \quad (4.13)$$

Setting  $1 - P(t) = s$  so that  $t = P^{-1}(1 - s)$ , we obtain  $\exp\left(\frac{t}{2\sigma^2}\right) \leq \frac{2N}{s}$ , which implies that  $t \leq 2\sigma^2 \log\left(\frac{2N}{s}\right)$ , and hence from (4.12) we get

$$\mathbb{E}[Z] \leq 2\sigma^2 \int_0^1 \log\left(\frac{2N}{s}\right) ds = 2\sigma^2 \log(2eN), \quad (4.14)$$

which is the thesis.  $\square$

To obtain a bound for  $\mathbb{E}[\|\mathbf{E}^n\|_\infty^2]^{1/2}$  we write each  $\mathbf{E}_i^n$  as

$$\mathbf{E}_i^n = \sum_{m=0}^{n-1} \sum_j (S(-\Delta t A)^{n-1-m})_{ij} \boldsymbol{\varepsilon}_j^m = \sum_{m=0}^{n-1} \sum_j Z_{m,j} = \sum_{k=1}^{(n-1)(K-1)^d} Z_k, \quad (4.15)$$

where we just re-indexed the last sum with  $k = \sum_{i=1}^d \mathbf{j}_i (K-1)^{i-1} + m(K-1)^d$ . Now, if we set  $P = (n-1)(K-1)^d$ ,  $X_p = \sum_{k=1}^P Z_k$ , with  $X_0 = 0$ , and  $X_P = \mathbf{E}_i^n$ , we have that  $\{X_p\}_{p=0}^P$  is a martingale: clearly  $\mathbb{E}[|X_p|] < \infty$  and

$$\mathbb{E}[X_p | X_0, \dots, X_{p-1}] = \mathbb{E}[Z_p | X_0, \dots, X_{p-1}] + X_{p-1} = X_{p-1}, \quad (4.16)$$

due to the properties of  $\boldsymbol{\varepsilon}^n$ , cf. Lemma 4.1. We can therefore apply Azuma-Hoeffding inequality with

$$|X_p - X_{p-1}| \leq \varepsilon |S(-\Delta t A)^{n-m-1}|_{ij} = c_p, \quad (4.17)$$

for  $p = \sum_{i=1}^d \mathbf{j}_i (K-1)^{i-1} + m(K-1)^d$ .  $\mathbf{E}_i^n$  is therefore sub-Gaussian with proxy variance

$$\sigma^2 = \sum_p c_p^2 = \varepsilon^2 \sum_{m=0}^{n-1} \sum_j ((S(-\Delta t A)^{n-m-1})_{ij})^2. \quad (4.18)$$

The next step is to derive a nicer expression for  $\sigma^2$ . We define the auxiliary vector  $\bar{\mathbf{E}}^n$ , which solves equation (4.1) in which  $\boldsymbol{\varepsilon}^n$  has been replaced by  $\bar{\boldsymbol{\varepsilon}}^n(\omega)$  with  $\bar{\boldsymbol{\varepsilon}}_i^n = \pm \varepsilon$  with equal probability and with  $\bar{\boldsymbol{\varepsilon}}_i^n$  and  $\bar{\boldsymbol{\varepsilon}}_j^n$  independent unless  $i = j$  and  $n = m$ . The key step now is to realize that by construction we have  $\sigma^2 = \mathbb{E}[(\bar{\mathbf{E}}_i^n)^2]$ , where the expectation can be bounded as follows.

$$\begin{aligned} \mathbb{E}[(\bar{\mathbf{E}}_i^n)^2] &= \varepsilon^2 \left(\frac{2}{K}\right)^{2d} \sum_{m=0}^{n-1} \sum_j \left( \sum_k S(s_k)^{n-1-m} (\mathbf{v}_k^h)_i (\mathbf{v}_k^h)_j \right)^2 \\ &= \varepsilon^2 \left(\frac{2}{K}\right)^{2d} \sum_{m=0}^{n-1} \sum_{\mathbf{k}, \bar{\mathbf{k}}} S(s_{\mathbf{k}})^{n-1-m} S(s_{\bar{\mathbf{k}}})^{n-1-m} (\mathbf{v}_{\mathbf{k}}^h)_i (\mathbf{v}_{\bar{\mathbf{k}}}^h)_i \sum_j (\mathbf{v}_{\mathbf{k}}^h)_j (\mathbf{v}_{\bar{\mathbf{k}}}^h)_j \\ &= \varepsilon^2 \left(\frac{2}{K}\right)^d \sum_{\mathbf{k}} \sum_{m=0}^{n-1} S(s_{\mathbf{k}})^{2(n-1-m)} (\mathbf{v}_{\mathbf{k}}^h)_i^2 \leq \varepsilon^2 \left(\frac{2}{K}\right)^d \sum_{\mathbf{k}} \frac{1}{1 - S(s_{\mathbf{k}})^2}. \end{aligned}$$

Here we used the fact that  $\sum_j (\mathbf{v}_k^h)_j (\mathbf{v}_{\tilde{k}}^h)_j = 0$  if  $k \neq \tilde{k}$  due to the orthogonality of the  $\mathbf{v}_k^h$  and the relations in (2.9). We can now apply Theorem 4.1 to obtain

$$\sigma = \mathbb{E}[(\tilde{\mathbf{E}}_i^n)^2]^{1/2} \leq 2^{d/2} \varepsilon K^{-d/2} \sqrt{\phi_d(\Delta t)} = 2^{d/2} \lambda^{-d/4} \varepsilon \Delta t^{d/4} \sqrt{\phi_d(\Delta t)}. \quad (4.19)$$

We conclude by invoking Theorem 4.5 and the fact that  $\mathbf{E}_i^n$  is sub-Gaussian with proxy variance  $\sigma^2$ , giving:

$$\mathbb{E}[\|\mathbf{E}^n\|_\infty^2]^{1/2} \leq 2^{(d+1)/2} \lambda^{-d/4} \varepsilon \Delta t^{d/4} \sqrt{\phi_d(\Delta t)} \sqrt{d/2 \log(\lambda^{-1} \Delta t) + \log(2e)}. \quad (4.20)$$

For fixed  $\lambda$ , the bound is  $O(\Delta t^{-1/4} |\log(\lambda^{-1} \Delta t)|^{1/2})$  in 1D,  $O(|\log(\lambda^{-1} \Delta t)|)$  in 2D and  $O(|\log(\lambda^{-1} \Delta t)|^{1/2})$  in 3D.

**REMARK 4.4** In this section we have shown that the use of SR leads to smaller error estimates. While these upper bounds<sup>2</sup> do not necessarily establish that SR will be more accurate than RtN, the fact that deterministic rounding modes lead to a linear growth of the error is well known (Henrici, 1962, 1963; Jézéquel, 1995; Connolly *et al.*, 2021). Indeed, we show in the next section that all our estimates are sharp, and that global errors behave as predicted by our bounds.

**REMARK 4.5** In this paper we did not discuss adaptive timestepping. We remark that this is not just for the sake of simplicity. Adaptivity in RK schemes typically relies on error estimates based on the comparison of two schemes that converge with different orders of accuracy. The problem when working in reduced precision is that these two schemes would not actually converge, but instead produce errors that, as we have shown in this section, stagnate or blow up as  $\Delta t \rightarrow 0$ . This behavior is likely to disrupt most adaptive schemes. The design of an adaptive strategy that is robust to rounding errors is currently an open question that we leave for future work.

## 5. Numerical results

In this section we present some numerical experiments that support and complete the theory presented in the paper. The code and data used to generate the figures and results of this section are publicly available on Zenodo<sup>3</sup>. All low-precision computations are emulated in software and performed using bfloat16. The fp16 format yields qualitatively similar results which we do not show here. Emulating low-precision operations is computationally expensive, especially for 2D and 3D problems. To make the computations faster, we employ a custom-built C++ emulator, Libchopping<sup>4</sup>, interfaced to Python, that combines vectorization, OpenMP and MPI for parallelization and uses the Intel Math Kernel Library for random number generation. The basic emulation algorithm was inspired by the low-precision emulator of Higham & Pranesh (2019). Other low-precision emulators supporting SR are also available (cf. Fousse *et al.* (2007); Zhang *et al.* (2019); Fasi & Mikaitis (2021b)), see Section 7(d) in Croci *et al.* (2022) for an overview.

**TEST PROBLEMS** Throughout the rest of the section, unless otherwise indicated, we fix  $h$  to be a power of two,  $\lambda = (1/2 - 2^{-4})/d$ ,  $T = \lceil 1/\Delta t \rceil \Delta t$  and  $g(\mathbf{x}) = G \in \mathbb{R}$  for all  $\mathbf{x}$ . The choice of  $\lambda$  was made so

<sup>2</sup>We remark that when working with rounding errors the only lower bound achievable is zero since floating-point computations can be exact in some specific scenarios.

<sup>3</sup>See <https://doi.org/10.5281/zenodo.6146386>.

<sup>4</sup>Available at <https://bitbucket.org/croci/libchopping/>.

as to satisfy the stability requirement for forward Euler (FE). Furthermore, we choose time-independent forcing terms so that the exact steady-state solutions are:

$$\begin{aligned} u_{1D}(\infty, x) &= (4x(1-x))^2 + G, & u_{2D}(\infty, \mathbf{x}) &= (16xy(1-x)(1-y))^2 + G, \\ u_{3D}(\infty, \mathbf{x}) &= (64xyz(1-x)(1-y)(1-z))^2 + G. \end{aligned}$$

Note that all of these solutions are symmetric with respect to  $x=\frac{1}{2}, y=\frac{1}{2}, z=\frac{1}{2}$ , satisfy the boundary conditions, and attain their maximum value of  $1+G$  at the center of the domain. To obtain the corresponding forcing terms it is sufficient to take the negative Laplacian of the exact solutions.

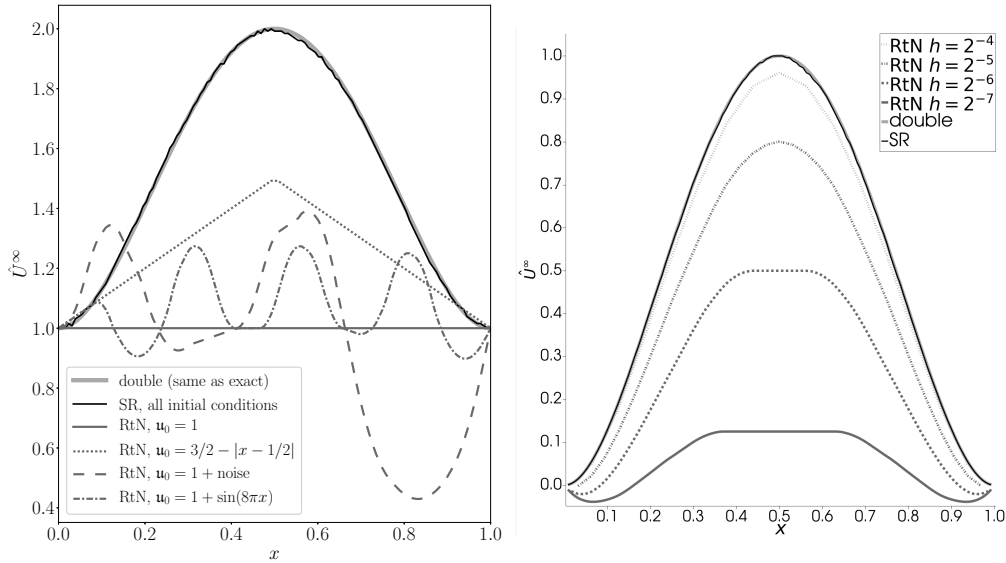


FIG. 1: Comparison between FE and exact steady-state solutions in 1D (left) and in 2D (right, here the plot is over the  $y = 1/2$  line in the  $xy$  plane) at  $t = T = 1$  for different initial conditions. In both plots we compare the exact steady-state solution (visibly the same as for double precision) with those obtained with `bfloat16+RtN` and `bfloat16+SR`. The SR solutions are always obtained with  $h = 2^{-7}$  and are as accurate as the precision allows. In the left figure, we fix  $G = 1$ ,  $h = 2^{-7}$  and we vary the initial condition leading to different steady-state RtN solutions. The noise term in the initial condition has been obtained by sampling independent standard Gaussian random variables at each mesh node. We note that all SR solutions converge to the same steady-state modulo some barely visible oscillations. In the right figure, we fix  $G = 0$ ,  $u_0 = 0$ , and, for RtN only, we vary the mesh size. The RtN solutions converge to smaller steady-state solutions as  $h$  is refined, significantly underestimating the true solution. On the finest grid the RtN graph even attains negative values.

In this section we mainly experiment with first order methods, namely FE in 1D, 2D and 3D and backward Euler (BE) in 1D and 2D. Experimentation with other popular Runge-Kutta methods, RK4, Crank-Nicolson, and first- and second-order RK-Chebyshev methods (RKC1 and RKC2; see [Verwer et al. \(1990\)](#)) yielded qualitatively similar results, which we mainly do not show here for the sake of brevity. The linear system solver used for BE was the Thomas algorithm in 1D and geometric multigrid

(V-cycles) with Jacobi relaxation in 2D. Both solvers were implemented in (emulated) low precision. For the geometric multigrid solver, we observe monotone convergence of the residual and we declare convergence when there is no longer any significant reduction in the residual. Experimentally, this only happens when the relative residual is close to machine precision.

Before proceeding, we show in Figure 1 what happens in practice when the heat equation is solved using double precision or bfloat16 with RtN and SR. The results in the left figure have been obtained with FE in 1D with  $h = 2^{-7}$  by setting  $g(x) = G = 1$  and by varying the initial condition. The results in the right figure have been obtained with FE in 2D by setting  $g(x) = G = 0$  and varying the mesh size:  $h \in \{2^{-l}\}_{l=4}^{l=7}$ . Both plots show the numerical and exact solutions at  $t = T = 1$ .

The results are in agreement with the considerations of Section 3.3. Due to stagnation, the low-precision RtN solution stops being updated and the values obtained significantly underestimate the true solution. In Figure 1 (left), we note how different initial conditions lead to different steady-state solutions for RtN. Conversely, the SR computations always converge to the same steady state (whenever this is representable) or mildly oscillate around it at random (when it is not representable). Remarkably, RtN solutions stagnate very close to low-frequency initial conditions ( $u_0 = 1$ ,  $u_0 = 3/2 - |x - 1/2|$ ) and later for high-frequency ones ( $u_0 = 1 + \sin(8\pi x)$  and Gaussian noise). In fact, the latter led to oscillatory solutions for which the discrete Laplacian, and consequently the  $\Delta \hat{\mathbf{U}}^n$  terms, are larger in magnitude and counterbalance the  $\Delta t$  in the stagnation condition (3.26).

In Figure 1 (right) we observe that stagnation occurs even for zero initial conditions. In this case, however, stagnation occurs gradually as the mesh size (hence  $\Delta t$ ) decreases, each time leading to different steady-state solutions. Furthermore, the steady-state solution seems to approach the zero initial condition as  $\Delta t \rightarrow 0$ . SR on the other hand is again impervious to stagnation and the bfloat16 results are as accurate as the working precision allows. In both plots we also observe that while RtN preserves symmetry due to its deterministic nature, SR does not. We note, however, that the asymmetry in the SR results is barely visible in the figure.

**Note:** From now on we set  $u_0 = G = 1$  in all the numerical experiments.

**REMARK 5.1** In our numerical experiments, the computation of the matrix-vector product  $-A\hat{\mathbf{U}}^n$  via two first-order differences (as described in Section 3.1) always resulted to be exact in all dimensions. This behavior is likely a consequence of  $h$  being a power of 2. Conversely, using a direct second-order difference introduced rounding errors growing like  $(h^{-2})$ , as predicted by Statement 1 of Theorem 3.2. We do not show the latter results here for the sake of brevity.

**LOCAL ROUNDING ERRORS** In Figure 2 we plot the worst-observed local rounding errors arising from the solution of the 2D test problem as a function of the timestep  $\Delta t$ . The 1D and 3D results, and the behavior of higher-order methods such as RK4 and RKC2 are qualitatively similar and we do not show them here. Recall that here  $\lambda$  is fixed and is the same for both FE and BE (we investigate how  $\lambda$  affects the BE results later). We consider three cases: 1) We use the delta form of the timestepping scheme and we use the more refined second-order differencing strategy (“delta form” in the figure). 2) We use the delta form, but we directly use a second order difference (“naive matvec” in the figure). 3) We do not use the delta form. We adopt the following error measures for the local rounding errors:

$$\frac{1}{u} \max_n \frac{\|\text{err}^n\|_\infty}{\|\hat{\mathbf{U}}^N\|_\infty}, \quad \text{for RtN}; \quad \frac{1}{u} \max_{n,\omega} \frac{\|\text{err}^n(\omega)\|_\infty}{\|\hat{\mathbf{U}}^N(\omega)\|_\infty}, \quad \text{for SR.} \quad (5.1)$$

Where  $\text{err}^n$  is either  $\Theta^n$  or  $\mathcal{E}^n$  depending on whether the delta form is used or not. For SR the max is taken over all observed realizations across the number of Monte Carlo samples used to estimate the

global errors (see later). We divide the errors by  $\|\hat{\mathbf{U}}^N\|_\infty$  to normalize the results: for this test problem the RtN solution has smaller magnitude due to stagnation, cf. Figure 1. We do not track the  $\boldsymbol{\varepsilon}^n$  term here since we know from Theorem 3.4 that it is bounded above by  $M\mathbf{v}$  where  $M = \max_n \|\hat{\mathbf{U}}^n\|_\infty$  giving a maximum value of 1 for RtN and of roughly 2 for SR for the error measures in (5.1). In fact, for this problem we observe that  $\|\hat{\mathbf{U}}^n\|_\infty \leq \|\hat{\mathbf{U}}^N\|_\infty$  for all  $n$  for RtN and that the same relation still approximately holds for SR.




FIG. 2: Worst-observed local errors as a function of  $\Delta t$  for RtN (left) and SR (right) and for FE (black) and BE (grey) in 2D. The slopes of the solid lines have been estimated from the observed values via linear regression. The results match with the theory showing that the delta form and the refined differencing strategy are significantly more accurate. Refining  $\Delta t$  further than the left figure values causes RtN to stagnate at  $\mathbf{U}^0 = \mathbf{1}$ .

From Figure 2 it is clear that using the delta form with the refined differencing strategy yields the minimum error. In this case Assumption (A3) appears to be satisfied since the error decreases with  $\Delta t$  at a rate  $p \in [1/2, 1]$ . For SR this rate eventually plateaus to  $p = 0$ . To understand better this behavior, we look at FE for simplicity and note that since the matrix-vector products are here computed exactly, we can write the (rounded) delta form update as

$$\Delta \hat{\mathbf{U}}^n + \tilde{\Theta}^n = (\Delta t A \hat{\mathbf{U}}^n + \Delta t (\mathbf{f}^n + \Delta \mathbf{f}^n))(1 + \boldsymbol{\delta}). \quad (5.2)$$

Here  $\Delta \mathbf{f}^n$  satisfies  $\|\Delta \mathbf{f}^n\|_\infty \leq C_f M_f \mathbf{v}$  (Assumption (A1) holds),  $\|\boldsymbol{\delta}\|_\infty < \mathbf{v}$  and we have ignored the error due to multiplication by  $\Delta t$  for simplicity. For small  $\Delta t$ , the  $\Delta t \Delta \mathbf{f}^n$  error term is negligible and the worst-case error in the update  $\Theta^n$  (cf. equation (3.3)) is, to leading order,

$$\Theta^n = O(\mathbf{v} \Delta t |A \hat{\mathbf{U}}^n|). \quad (5.3)$$

Let  $m \in \{0, 1, 2\}$ . If  $\widehat{\mathbf{U}}^n$  were the discrete version of an  $m$ -times differentiable function of space, we would expect  $\Delta t |A\widehat{\mathbf{U}}^n| = O(\Delta t h^{m-2}) = O(\Delta t^m)$ . Comparing this rate with Figure 2, it then appears that both RtN and SR solution are less smooth than the exact solution  $u$ . Furthermore, the SR solution smoothness deteriorates as the timestep is reduced.<sup>5</sup> Note that the RtN rate here does not deteriorate since for small  $\Delta t$  stagnation occurs at the initial condition (cf. Section 3.3), which is smooth and exactly representable.

Comparing Figure 2 with what is stated in Theorem 3.4 we note that these experiments mostly match the theoretical predictions. We observe a  $\theta$  of order  $O(\Delta t^p)$  with  $p > 1/2$  (at least initially) when the refined differencing strategy is used and  $p = 0$  when it is not. Furthermore, all local errors are essentially bounded as  $\Delta t$  is reduced. Finally, these results also show that when the refined differencing strategy is used  $\Theta^n$  is negligible with respect to  $\epsilon^n$ , which is  $O(1)$  in the worst-case. We therefore expect the  $\epsilon^n$  terms to dominate the global error.

**GLOBAL ROUNDING ERRORS** We conclude the section by analysing the global rounding errors. From now on we only work with the delta form and the refined differencing strategy since these clearly give the most accurate implementation. We just saw how in this case the  $\epsilon^n$  term dominates over  $\Theta^n$  so we expect the global errors to behave according to Scenario 1) (cf. Section 4) for RtN and to Scenario 2) for SR. We monitor the following quantities:

$$u^{-1} \frac{\|\mathbf{E}^N\|}{\|\widehat{\mathbf{U}}^N\|}, \text{ for RtN}; \quad u^{-1} \frac{\mathbb{E}[\|\mathbf{E}^N(\omega)\|^2]^{1/2}}{\|\mathbf{U}^N\|}, \text{ for SR}, \quad (5.4)$$

where the norms used are either the infinity or the discrete  $L^2$  norm. We use standard Monte Carlo sampling to estimate the expected errors in the SR case. Rather than using a fixed number of samples, we use a number which ensures that the true expected error is within 5% of the estimated value with at least 95% confidence. We work with both FE and BE, as well as the RK4 and RKC2 (16 stages) methods, which we all compare by using the same  $\Delta t$  values for  $\lambda$  fixed. We remark that this is not the standard way of employing RKC methods<sup>6</sup>. Nevertheless, we believe this is a good way of comparing RKC2 with the other methods in the asymptotic regime  $\Delta t \rightarrow 0$ .

**REMARK 5.2** In all practical simulations with RtN, the error measures (5.4) always satisfied

$$u^{-1} \frac{\|\text{fl}(\mathbf{U}^\infty) - \mathbf{U}^\infty\|}{\|\mathbf{U}^\infty\|} \leq \text{error measure} \leq u^{-1} \frac{\|\widehat{\mathbf{U}}^0 - \mathbf{U}^\infty\|}{\|\mathbf{U}^\infty\|}, \quad (5.5)$$

where  $\widehat{\mathbf{U}}^0 = \mathbf{U}_i^0 = u(0, \mathbf{x}_i)$ . The lower bound is actually a theoretical result that also applies to SR (after replacing  $\text{fl}(\cdot)$  with  $\text{sr}(\cdot)$  and taking the minimum over  $\omega$ ). In fact, it indicates that the best we can hope for is to compute the discrete steady-state solution exactly and then rounding it. On the other hand, the upper bound confirms (cf. Section 3.3) that in the worst case RtN stagnation begins from the very first timestep so that the rounded solution is stuck at the initial condition. Note that in (5.4) we divide the RtN error by  $\|\widehat{\mathbf{U}}^N\|$  rather than  $\|\mathbf{U}^N\|$ , to make the comparison with SR fairer when stagnation occurs.

<sup>5</sup>For a visible example of what happens in practice, observe the oscillations in the SR solutions shown in Figure 1

<sup>6</sup>The stability region of RKC methods grows quadratically with the number of stages, which are increased as needed when the grid size is refined (Verwer *et al.*, 1990). Therefore, the common usage of RKC2 would be to increase the number of stages and  $\lambda$  proportionally to  $1/h$  for  $\Delta t$  fixed.



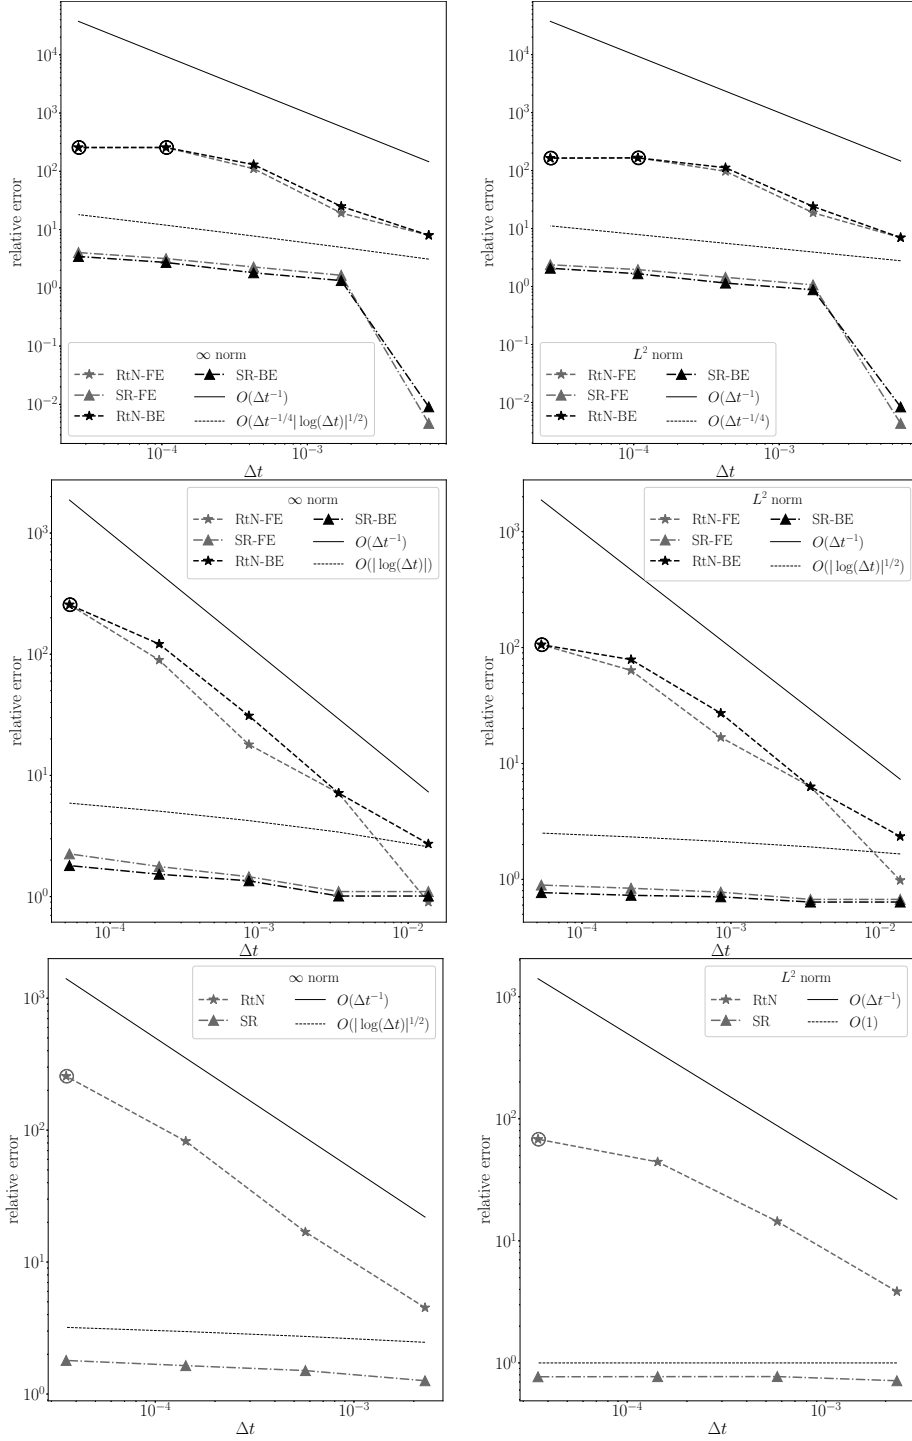


FIG. 3: Plots of the relative global rounding error in the  $\infty$  (left) and discrete  $L^2$  norm (right) in 1D (top), 2D (middle), and 3D (bottom). We circled the RtN data points for which the solution stagnates at the initial condition. The error behavior matches the theoretical predictions of Section 4. On the coarsest mesh the 1D steady-state solution and forcing term is (almost) exactly representable, hence the low error SR data points in the 1D plots.

Results are shown in Figures 3 and 4. The SR results essentially follow the asymptotic rates derived in Section 4 and described in Table 2, with possibly slightly lower rates in 1D (cf. Figure 3 top). These results show that SR behaves as predicted by our theory and that the  $\epsilon^n$  term is indeed dominating the local error, thus determining the global error behavior. Comparing the SR results with the lower bound in (5.5), we note that the  $O(1)$  errors observed with SR are extremely close to the best we can theoretically hope to achieve. On the other hand, when RtN is used the relative global errors are always growing like  $O(\Delta t^{-1})$  until severe stagnation occurs, matching the theoretical bounds derived in Section 4. This confirms that in the RtN case rounding errors cannot be modelled as spatially or temporally uncorrelated and/or independent: if this were the case a smaller growth rate would be observed due to cancellation of cross-correlation terms in the global error analysis of Section 4.

We observe that when severe stagnation occurs our RtN global error analysis becomes too pessimistic. In fact, the local errors stop accumulating in the worst possible way as the RtN solution stagnates at the initial condition (cf. Remark 5.2). Nevertheless, our bounds are sharp when they are needed, as the severe stagnation regime should just be avoided. We remark that for neither rounding modes we observe any additional growth in  $\Delta t$  due to implicit timestepping, indicating that in this case the constants appearing in Theorem 3.4 might be  $K$ -independent. In fact, the FE and BE global rounding errors are comparable in size.

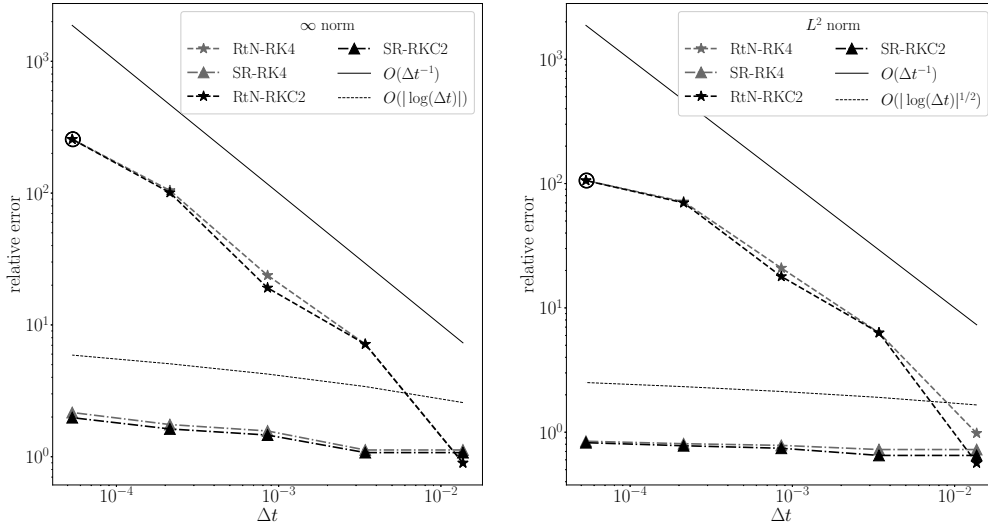


FIG. 4: Plots of the relative global rounding error in the  $\infty$  (left) and discrete  $L^2$  norm (right) in 2D for the RK4 and second-order RK-Chebyshev (16 stages) methods. We circled the RtN data points for which the solution stagnates at the initial condition. The global error behavior matches the theoretical predictions of Section 4. Comparing these results with Figure 3 (middle), these higher-order methods do not seem to improve over first-order methods.

In Figure 4 we specifically look at the global error behavior of the higher-order methods RK4 and RKC2 in 2D. Interestingly, the error behavior of high-order methods is qualitatively the same as for the first-order Euler methods. In fact, our theory holds for all RK methods, and states that no matter the order, RK methods will stop converging in reduced-precision with the error either stagnating to an  $O(u)$  value (for SR in 3D) or blow-up as  $\Delta t \rightarrow 0$  (with a much faster error growth rate for RtN). With these results in mind, we suspect that there might not actually be any accuracy and performance benefit

of using higher-order methods in reduced precision, although they might still be required for nonlinear stability purposes.

REMARK 5.3 The SR error lines in Figure 3 all behave according to Scenario 2) (cf. beginning of Section 4). This suggests that the  $\theta$  term behaves better than the worst case scenario (cf. Remark 4.1) and that the  $\varepsilon$  term dominates. We suspect that the  $\theta$  term behaves somewhere in between the two scenarios, with local rounding errors that are approximately uncorrelated (or mean-independent) across timesteps. We leave this investigation to future research.

Finally, in Figure 5 we show the behavior of the global rounding error as a function of  $\lambda = \{2^{(l+5)} + 5\}_{l=0}^7$ , for fixed  $h = 2^{-7}$  with BE in 1D. The motivation behind this experiment is that we expect the errors to worsen as the condition number of the BE linear system (which is  $O(\lambda)$  in this case) grows, as mentioned in Remark 3.3. Indeed this is the case: when  $\lambda \approx u^{-1}$  the rounding error starts increasing dramatically. This behavior worsens in higher dimensions where the condition number is larger and more operations are needed by the linear solver (not shown). These results suggest that  $\lambda$ , and consequently  $\Delta t$ , must be kept small even for implicit timestepping, thus restricting what is perhaps the greatest advantage of implicit methods. We suggest that a mixed-precision treatment (Abdelfattah *et al.*, 2021) would be extremely beneficial here.




FIG. 5: Plot of the relative global rounding error in the  $\infty$  norm for BE in 1D as a function of  $\lambda$  for fixed  $h = 2^{-7}$ . The global error increases drastically as the conditioning of the BE linear system increases beyond roughly  $u^{-1}$ .

## 6. Conclusions

We have investigated the behavior of local and global rounding errors arising from the solution of the heat equation in low precision using both round-to-nearest and stochastic rounding. We showed that a careful implementation of the scheme in delta form is essential to reduce rounding errors, and we explained why round-to-nearest computations are always affected by stagnation as  $\Delta t$  decreases. For this rounding mode, the global rounding error grows like  $O(u\Delta t^{-1})$  until stagnation occurs and the discrete solution converges to the initial condition for  $\Delta t$  small enough. On the other hand, we showed that stochastic rounding successfully prevents stagnation and always approximates the exact solution. Furthermore, the dominant local error terms arising from stochastic rounding have favorable statistical properties, yielding a much milder global error growth rate, which is  $O(u\Delta t^{-1/4})$  in 1D and roughly constant in higher dimensions. These results suggest that there might be little accuracy and performance

benefit with using high-order RK schemes in reduced precision, and that a mixed-precision treatment could be more suitable for these methods.

We remark that the global error behavior described in this paper occurs regardless of the number of digits used. However, rounding errors of this type only grow appreciably when working in low precision. In this setting, our results show that stochastic rounding empowers the PDE solver with resilience to rounding errors, which enables the computation of (almost) machine-precision-accurate solutions.

While our theory and error bounds broadly also extend to linear parabolic PDEs in general (a paper on this subject is in preparation), it would be extremely interesting to analyse the effect of SR when solving other time-dependent PDEs, especially hyperbolic or nonlinear problems. In this setting, stagnation might not necessarily occur, and SR might consequently not be as effective. We leave this investigation to future work.

## REFERENCES

- ABDELFATTAH, A., ANZT, H., BOMAN, E. G., CARSON, E., COJEAN, T., DONGARRA, J., FOX, A., GATES, M., HIGHAM, N. J., LI, X. S. *et al.* (2021) A survey of numerical linear algebra methods utilizing mixed-precision arithmetic. *The International Journal of High Performance Computing Applications*, **35**, 344–369.
- ALBEN, J., MICIKEVICIUS, P., WU, H. & SIU, M. Y., NVIDIA Corporation (12 Dec 2019) Stochastic rounding of numerical values. Patent, US 2019/0377549 A1.
- ARATÓ, M. (1983) Round-off error propagation in the integration of ordinary differential equations by one step methods. *Acta Scientiarum Mathematicarum*, **45**, 23–31.
- AZUMA, K. (1967) Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal*, **19**.
- BRADBURY, J. D., CARLOUGH, S. R., PRASKY, B. R. & SCHWARZ, E. M., International Business Machines Corporation (3 Jul 2017b) Stochastic rounding floating-point multiply instruction using entropy from a register. Patent, US 2017/0220341 A1.
- BRADBURY, J. D., CARLOUGH, S. R., PRASKY, B. R. & SCHWARZ, E. M., International Business Machines Corporation (6 Jul 2017a) Reproducible stochastic rounding for out of order processors. Patent, US 2017/0192749 A1.
- BUTCHER, J. C. (2016) *Numerical Methods for Ordinary Differential Equations*, 3 edn. John Wiley & Sons, Ltd.
- CONNOLLY, M. P., HIGHAM, N. J. & MARY, T. (2021) Stochastic rounding and its probabilistic backward error analysis. *SIAM Journal on Scientific Computing*, **43**, A566–A585.
- COURANT, R. & HILBERT, D. (1953) *Methods of Mathematical Physics*, vol. 1. New York: Interscience Publishers: Wiley Classic Library.
- CROCI, M., FASI, M., HIGHAM, N. J., MARY, T. & MIKAITIS, M. (2022) Stochastic rounding: Implementation, error analysis, and applications. *Royal Society Open Science* (to appear).
- DAVIES, M., SRINIVASA, N., LIN, T. H., CHINYA, G., CAO, Y., CHODAY, S. H., DIMOU, G., JOSHI, P., IMAM, N., JAIN, S., LIAO, Y., LIN, C. K., LINES, A., LIU, R., MATHAIKUTTY, D., MCCOY, S., PAUL, A., TSE, J., VENKATARAMANAN, G., WENG, Y. H., WILD, A., YANG, Y. & WANG, H. (2018) Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro*, **38**, 82–99.
- DÜBEN, P. D., SUBRAMANIAN, A., DAWSON, A. & PALMER, T. (2017) A study of reduced numerical precision to make superparameterization more competitive using a hardware emulator in the OpenIFS model. *Journal of Advances in Modeling Earth Systems*, **9**, 566–584.
- ESSAM, M., TANG, T. B., HO, E. T. W. & CHEN, H. (2017) Dynamic point stochastic rounding algorithm for limited precision arithmetic in deep belief network training. *International IEEE/EMBS Conference on Neural Engineering (NER)*. Shanghai: IEEE, pp. 629–632.
- FASI, M. & MIKAITIS, M. (2021a) Algorithms for stochastically rounded elementary arithmetic operations in IEEE 754 floating-point arithmetic. *IEEE Transactions on Emerging Topics in Computing*.

- FASI, M. & MIKAITIS, M. (2021b) Algorithms for stochastically rounded elementary arithmetic operations in ieee 754 floating-point arithmetic. *IEEE Transactions on Emerging Topics in Computing*, **9**, 1451–1466.
- FORSYTHE, G. E. (1950) Round-off errors in numerical integration on automatic machinery. Preliminary report. *Bulletin of the American Mathematical Society*, **56**, 61–61.
- FORSYTHE, G. E. (1959) Reprint of a note on rounding-off errors. *SIAM Review*, **1**, 66–67.
- FOUSSE, L., HANROT, G., LEFÈVRE, V., PÉLISSIER, P. & ZIMMERMANN, P. (2007) Mpr: A multiple-precision binary floating-point library with correct rounding. *ACM Transactions on Mathematical Software (TOMS)*, **33**, 13–es.
- GUPTA, S., AGRAWAL, A., GOPALAKRISHNAN, K. & NARAYANAN, P. (2015) Deep learning with limited numerical precision. *Proceedings of the 32nd International Conference on Machine Learning*. Lille, France: W&CP.
- HAIRER, E. & WANNER, G. (1996) *Solving Ordinary Differential Equations II*, vol. 14, 2 edn. Springer-Verlag Berlin Heidelberg, pp. 1–614.
- HENRICI, P. (1962) *Discrete Variable Methods in Ordinary Differential Equations*, 1 edn. John Wiley & Sons, Inc., pp. 1–407.
- HENRICI, P. (1963) *Error Propagation in Difference Methods*, 1 edn. John Wiley & Sons, Inc., pp. 1–73.
- HENRY, G. G. & REED, D. R., Alliance Semiconductor Co., Ltd. (7 Jun 2018) Processor with memory array operable as either cache memory or neural network unit memory. Patent, US 2018/0157970 A1.
- HIGHAM, N. J. (2002) *Accuracy and Stability of Numerical Algorithms*. SIAM.
- HIGHAM, N. J. (2008) *Functions of Matrices: Theory and Computation*. SIAM.
- HIGHAM, N. J. & MARY, T. (2019) A new approach to probabilistic rounding error analysis. *SIAM Journal of Scientific Computing*, **41**, 2815–2835.
- HIGHAM, N. J. & MARY, T. (2020) Sharper probabilistic backward error analysis for basic linear algebra kernels with random data. *SIAM Journal on Scientific Computing*, **42**, A3427–A3446.
- HIGHAM, N. J. & PRANESH, S. (2019) Simulating low precision floating-point arithmetic. *SIAM Journal on Scientific Computing*, **41**, C585–C602.
- HOEFFDING, W. (1963) Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, **58**, 13–30.
- HOPKINS, M., MIKAITIS, M., LESTER, D. R. & FURBER, S. (2020) Stochastic rounding and reduced-precision fixed-point arithmetic for solving neural ordinary differential equations. *Philosophical Transactions of the Royal Society A*, **378**, 20190052.
- HULL, T. E. & SWENSON, J. R. (1966) Tests of probabilistic models for propagation of roundoff errors. *Communications of the ACM*, **9**, 108–113.
- IEEE 754 (2008) IEEE Standard for Floating-Point Arithmetic, IEEE Std 754-2008 (revision of IEEE Std 754-1985). <https://doi.org/10.1109/IEEESTD.2008>.
- IPSEN, I. C. & ZHOU, H. (2020) Probabilistic error analysis for inner products. *SIAM Journal on Matrix Analysis and Applications*, **41**, 1726–1741.
- JÉZÉQUEL, F. (1995) Round-off error propagation in the solution of the heat equation by finite differences. *Journal of Universal Computer Science*, **1**, 465–479.
- JIA, Z., TILLMAN, B., MAGGIONI, M. & SCARPAZZA, D. P. (2019) Dissecting the Graphcore IPU architecture via microbenchmarking. *Technical Report*. Chicago: Citadel High Performance Computing R&D Team. <https://arxiv.org/pdf/1912.03413>.
- KANTER, O. A. & BAR, I., Densbits Technologies Ltd. (27 May 2010) Apparatus and methods for hardware-efficient unbiased rounding. Patent, US 2010/0131580 A1.
- KAUL, H., ANDERS, M. A., MATHEW, S. K., YAO, A., RAY, J., TANG, P. T., STRICKLAND, M. S., CHEN, X., SHPEISMAN, T., R., A. A., KOKER, A., SINHA, K., VEMBU, B., GALOPPO VON BORRIES, N. C., NURVITADHI, E., BARIK, R., LIN, T., RANGANATHAN, V. & JAHAGIRDAR, S., Intel Corporation (1 Nov 2018) Instructions and logic to perform floating-point and integer operations for machine learning. Patent, US

- 2018/0315398 A1.
- KLOEDEN, P. E. & PLATEN, E. (2013) *Numerical Solution of Stochastic Differential Equations*, vol. 23. Springer Science & Business Media.
- KLÖWER, M., HATFIELD, S., CROCI, M., DÜBEN, P. & PALMER, T. (2021) Fluid simulations accelerated with 16 bit: Approaching 4x speedup on a64fx by squeezing shallowwaters.jl into float16. *Earth and Space Science Open Archive*, 26. <https://doi.org/10.1002/essoar.10507472.2>.
- KUTTLER, J. R. (1974) Direct methods for computing eigenvalues of the finite-difference Laplacian. *SIAM Journal of Numerical Analysis*, **11**.
- LIFSCHES, S., GSI Technology Inc. (9 Jan 2020) In-memory stochastic rounder. Patent, US 2020/0012708 A1.
- LOH, G. H., Advanced Micro Devices, Inc. (26 Sep 2019) Stochastic rounding logic. Patent, US 2019/0294412 A1.
- LOMAX, H., PULLIAM, T. & ZINGG, D. (2013) *Fundamentals of Computational Fluid Dynamics*. Scientific Computation. Springer Berlin Heidelberg.
- MELLEMPUDI, N., SRINIVASAN, S., DAS, D. & KAUL, B. (2019) Mixed precision training with 8-bit floating point. *Technical Report*. <https://arxiv.org/abs/1905.12334v1>.
- MIKAITIS, M. (2021) Stochastic rounding: Algorithms and hardware accelerator. *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, pp. 1–6.
- NA, T., KO, J. H., KUNG, J. & MUKHOPADHYAY, S. (2017) On-chip training of recurrent neural networks with limited numerical precision. *Proceedings of the International Joint Conference on Neural Networks*. Anchorage, AK, USA: IEEE, pp. 3716–3723.
- ORTIZ, M., CRISTAL, A., AYGUADÉ, E. & CASAS, M. (2018) Low-precision floating-point schemes for neural network training. *Technical Report*. <https://arxiv.org/pdf/1804.05267>.
- PAXTON, E. A., CHANTRY, M., KLÖWER, M., SAFFIN, L. & PALMER, T. (2021) Climate modelling in low-precision: effects of both deterministic & stochastic rounding. *Technical Report*. <https://arxiv.org/pdf/2104.15076>.
- SEBASTIAN, H. & MAYR, C. (2018) SpiNNaker2-Towards extremely efficient digital neuromorphics and multi-scale brain emulation. *2018 NICE Workshop*. NICE. <https://niceworkshop.org/%0Awp-content/uploads/2018/05/27-SHoppner-SpiNNaker2.pdf>.
- VÁÑA, F., DÜBEN, P., LANG, S., PALMER, T., LEUTBECHER, M., SALMOND, D. & CARVER, G. (2017) Single precision in weather forecasting models: An evaluation with the IFS. *Monthly Weather Review*, **145**, 495–502.
- VERWER, J. G., HUNDSDORFER, W. H. & SOMMEIJER, B. P. (1990) Convergence properties of the Runge-Kutta-Chebyshev method. *Numerische Mathematik*, **57**, 157–178.
- VON NEUMANN, J. & GOLDSTINE, H. H. (1947) Numerical inverting of matrices of high order. *Bulletin of the American Mathematical Society*, **53**, 1021–1099.
- WANG, N., CHOI, J., BRAND, D., CHEN, C.-Y. & GOPALAKRISHNAN, K. (2018) Training deep neural networks with 8-bit floating point numbers. *32nd Conference on Neural Information Processing Systems (NeurIPS 2018)*. Montréal, Canada: NIPS, pp. 1–10.
- ZHANG, T., LIN, Z., YANG, G. & DE SA, C. (2019) Qpytorch: A low-precision arithmetic simulation framework. *2019 Fifth Workshop on Energy Efficient Machine Learning and Cognitive Computing-NeurIPS Edition (EMC2-NIPS)*. IEEE, pp. 10–13.

### Appendix A. Proof of Theorem 3.3

First, we need an auxiliary lemma:

**LEMMA A.1** Let  $P(z)$  be a polynomial of degree  $s$ , let  $\mathbf{v} \in \mathbb{R}^{(K-1)^d}$  and set  $Y = P(-\Delta t A)$  and  $\mathbf{y} = P(-\Delta t A)\mathbf{v}$ . There exist constants  $c_0, c_1, c_2 > 0$  such that  $\|\mathbf{y}\|_\infty \leq c_0(s, d, \lambda) \|\mathbf{v}\|_\infty$  and, for  $\mathbf{v}$  sufficiently small, the computation of  $Y$  and  $\mathbf{y}$  in finite precision yields instead  $\hat{Y}$  and  $\hat{\mathbf{y}}$  satisfying

$$\|\hat{Y} - Y\|_\infty \leq c_1(s, d, \lambda) \mathbf{v}, \quad \|\hat{\mathbf{y}} - \mathbf{y}\|_\infty \leq c_2(s, d, \lambda) \mathbf{v} \|\mathbf{v}\|_\infty. \quad (\text{A.1})$$

The second bound still holds, albeit with a slightly larger constant, in the case in which  $\mathbf{v}$  is also computed inexactly yielding  $\hat{\mathbf{v}}$  with  $\|\hat{\mathbf{v}} - \mathbf{v}\|_\infty \leq c \mathbf{v} \|\mathbf{v}\|_\infty$  for some  $c > 0$ .

*Proof.* Let  $P(x) = \sum_{q=0}^s a_q x^q$  and, for simplicity of notation, let  $B = -\Delta t A$ , where  $B$  is sparse with at most  $m = 3d$  nonzero entries per row. From (2.5) it is easy to see that  $\|B\|_\infty = 4d\lambda$ . The first bound is obtained by noting that since  $\mathbf{y} = P(B)\mathbf{v}$  we have

$$\|\mathbf{y}\|_\infty \leq \sum_{q=0}^s |a_q| \|B\|_\infty^q \|\mathbf{v}\|_\infty \leq c_0 \|\mathbf{v}\|_\infty, \quad \text{where} \quad c_0 = \max_q |a_q| \frac{(4d\lambda)^{s+1} - 1}{4d\lambda - 1}.$$

Now, if the evaluation of  $P(B)$  is performed via explicit powers, Horner's method or the Paterson-Stockmeyer method (cf. Section 4.2 in Higham (2008)), then Theorem 4.5 in Higham (2008) (see Section 5 of Higham (2002) for a proof) gives, for some  $\bar{c} > 0$  depending on  $s$  and  $m$ ,

$$\|\hat{Y} - Y\|_\infty \leq \bar{c} \mathbf{v} \sum_{q=0}^s |a_q| \|B\|_\infty^q \leq c_1(s, d, \lambda) \mathbf{v},$$

where  $c_1(s, d, \lambda) = \bar{c} c_0$ . Note that  $\bar{c}$  here does not depend on the size of  $B$  since  $B$  is sparse. When  $\mathbf{v}$  is computed exactly, the other bound in (A.1) is obtained with the same argument by using the classical error bound for sparse matrix-vector multiplication (see Section 2 in Higham (2002)):

$$\widehat{B}\mathbf{v} = (B + \Delta B)\mathbf{v}, \quad \text{with} \quad |\Delta B| \leq \gamma_m |B|.$$

When  $\mathbf{v}$  is evaluated inexactly we can apply the bound for the exact vector evaluation to  $\hat{\mathbf{v}}$ , giving

$$\|\hat{\mathbf{y}} - \mathbf{y}\|_\infty \leq c_2 \mathbf{v} \|\hat{\mathbf{v}}\|_\infty \leq c_2 \mathbf{v} (1 + c \mathbf{v}) \|\mathbf{v}\|_\infty,$$

and the proposition follows after renaming the constant in the bound.  $\square$

We can now prove Theorem 3.3.

*Proof of Theorem 3.3.* Lemma A.1 already gives the result for explicit RK methods. For implicit methods, let  $B = -\Delta t A$ . We evaluate  $\mathbf{z} = P(B)\mathbf{v}$  by solving the linear system  $W(B)\mathbf{z} = N(B)\mathbf{v} = \mathbf{y}$ . Since  $W(x)$  has no roots on the negative real axis and  $B$  only has negative eigenvalues,  $W(B)$  is invertible and there exists  $\alpha > 0$  such that  $\|W(B)^{-1}\|_\infty \leq \alpha(s, d, \lambda, K)$ . Thanks to Lemma A.1 we then obtain  $\|\mathbf{z}\|_\infty \leq \alpha c_0 \|\mathbf{v}\|_\infty$ , which is the first preposition. Now, due to Lemma A.1, the evaluation of  $\mathbf{z}$  in finite precision yields

$$(W(B) + \Delta W)\hat{\mathbf{z}} = \mathbf{y} + \Delta \mathbf{y}, \quad \|\Delta W\|_\infty \leq c_1 \mathbf{v}, \quad \|\Delta \mathbf{y}\|_\infty \leq c_2 \mathbf{v} \|\mathbf{v}\|_\infty$$

where  $\hat{\mathbf{z}}$  comes from the finite-precision solution of the perturbed linear system. Since the system is solved by a stable method by assumption, we can combine the first preposition of the theorem with Theorem 7.2 in Higham (2002) to obtain, for some constant  $\bar{c} > 0$ ,

$$\|\hat{\mathbf{z}} - \mathbf{z}\|_\infty \leq \bar{c}(s, d, \lambda, K) \mathbf{v} (c_1 \alpha + c_2 \alpha^2 c_0) \|\mathbf{v}\|_\infty = C(s, d, \lambda, K) \mathbf{v} \|\mathbf{v}\|_\infty.$$



This is the proposition that we wanted to prove. When  $\mathbf{v}$  is evaluated inexactly we can apply the bound for the exact vector evaluation to  $\hat{\mathbf{v}}$  and then use the triangle inequality in the same way as in the proof of Lemma A.1.  $\square$

## Appendix B. Proof of Theorem 4.1

We first need two auxiliary results:

LEMMA A.1 The  $s_{\mathbf{k}}$  and  $\mathbf{v}_{\mathbf{k}}^h$  all satisfy

$$-4d\lambda \leq s_{\mathbf{k}} \leq -d\pi^2 \Delta t \left( 1 - \frac{\pi^2}{12} \lambda^{-1} \Delta t \right) \leq 0, \quad (\text{A.1})$$

$$-\Delta t \pi^2 \|\mathbf{k}\|_2^2 \leq s_{\mathbf{k}} \leq -\left( 1 - \frac{\pi^2}{12} \right) \Delta t \pi^2 \|\mathbf{k}\|_2^2, \quad (\text{A.2})$$

$$0 \leq (\mathbf{1}, \mathbf{v}_{\mathbf{k}}^h) = \prod_{j=1}^d \cot \left( \frac{\pi \mathbf{k}_j}{2K} \right) (\mathbf{k}_j \bmod 2) \leq 2^d \pi^{-d} K^d \prod_{j=1}^d \frac{1}{\mathbf{k}_j} (\mathbf{k}_j \bmod 2). \quad (\text{A.3})$$

where the upper and lower bound of the first inequality are only attained in the limit as  $\Delta t, h \rightarrow 0$  and strict inequalities hold otherwise.

*Proof.* The relations involving the  $s_{\mathbf{k}}$  terms are obtained by using the inequality  $x^2 - x^4/3 \leq \sin(x)^2 \leq \min(1, x^2)$ , valid for all  $x \in \mathbb{R}$ , to (2.8). The definition of  $\mathbf{v}_{\mathbf{k}}^h$  in (2.8) gives

$$(\mathbf{1}, \mathbf{v}_{\mathbf{k}}^h) = \prod_{j=1}^d \left( \sum_{i_j=1}^{K-1} \sin \left( \pi i_j \frac{\mathbf{k}_j}{K} \right) \right). \quad (\text{A.4})$$

Applying Lagrange trigonometric identity to the series of sines yields the equality in (A.3). The upper bound is obtained by using the relation  $\cot(x) \leq 1/x$ , valid for all  $x \in (0, \pi)$ .  $\square$

LEMMA A.2 The following bounds hold for  $K \geq 1$  and  $\mathbf{k} \in \mathbb{N}^d$ :

$$\sum_{\mathbf{k}=1}^{K-1} \frac{1}{\|\mathbf{k}\|_2^2} \leq \begin{cases} \frac{\pi^2}{6}, & d=1, \\ \frac{\pi}{2} \log(\sqrt{2}K) + \frac{1}{2}, & d=2, \\ \frac{\pi\sqrt{3}}{2}K - \frac{\pi}{2} + \frac{1}{3}, & d=3, \end{cases} \quad (\text{A.5})$$

$$\sum_{\mathbf{k}=1}^{K-1} \frac{1}{\|\mathbf{k}\|_2^4} \leq \begin{cases} \frac{\pi^4}{90}, & d=1, \\ \frac{\pi}{4} + \frac{1}{4}, & d=2, \\ \frac{\pi}{2} + \frac{1}{9}, & d=3. \end{cases} \quad (\text{A.6})$$

*Proof.* The 1D result is actually an equality and is well-known. Let  $r = \|\mathbf{x}\|_2$ . If  $f(r)$  is strictly positive and monotonically decreasing on  $[1, \infty)$ , then, in 2D

$$\begin{aligned} f(\sqrt{i^2 + j^2}) &< \int_{i-1}^i \int_{j-1}^j f(\sqrt{x^2 + y^2}) dx dy, \\ \Rightarrow \sum_{\mathbf{k}=1}^{K-1} f(\|\mathbf{k}\|) &\leq f(\sqrt{2}) + \int_1^{\sqrt{2}K} \frac{\pi}{2} r f(r) dr, \end{aligned}$$

while in 3D

$$\begin{aligned} f(\sqrt{i^2 + j^2 + k^2}) &< \int_{i-1}^i \int_{j-1}^j \int_{k-1}^k f(\sqrt{x^2 + y^2 + z^2}) dx dy dz, \\ \Rightarrow \sum_{k=1}^{K1} f(\|k\|) &\leq f(\sqrt{3}) + \int_1^{\sqrt{3}K} \frac{\pi}{2} r^2 f(r) dr. \end{aligned}$$

The desired bounds then follow naturally using  $f(r) \equiv r^{-2}$  for (A.5) and  $f(r) \equiv r^{-4}$  for (A.6).  $\square$

Finally, we can prove Theorem 4.1:

*Proof.* We start by proving (4.5). Since  $(1 - x^2)^{-1} \leq (1 - |x|)^{-1}$  for all  $x$  we just need to prove the second inequality in (4.5). Since  $S(z) = 1 + z + O(z^2)$  and  $(\Lambda \cup \{-4d\lambda\}) \subset \mathcal{A}$ , there exist  $s_0, \alpha_0 > 0$  such that

$$0 \leq S(z) \leq 1 + \frac{z}{2}, \text{ for } z \in (-s_0, 0), \text{ and } \alpha_0 \leq 1 - |S(z)| \leq 1, \text{ for } z \in [-4d\lambda, -s_0].$$

We can therefore split

$$\sum_k \frac{1}{1 - |S(s_k)|} = \sum_{|s_k| \leq s_0} \frac{1}{1 - |S(s_k)|} + \sum_{|s_k| > s_0} \frac{1}{1 - |S(s_k)|} \leq \sum_{|s_k| \leq s_0} \frac{2}{-s_k} + \frac{(K-1)^d}{\alpha_0}, \quad (\text{A.7})$$

Due to (A.2), and since  $-s_k > 0$  for all  $k$ , we now have

$$\sum_{|s_k| \leq s_0} \frac{1}{-s_k} \leq \frac{1}{c\Delta t} \sum_k \frac{1}{\|k\|_2^2}, \quad \text{where } c = \pi^2 \left(1 - \frac{\pi^2}{12}\right). \quad (\text{A.8})$$

By combining the above with (A.7) and Lemma A.2 we obtain (4.5) since

$$\sum_k \frac{1}{1 - |S(s_k)|} \leq \frac{2\Delta t^{-1}}{c} \sum_k \frac{1}{\|k\|_2^2} + \frac{(K-1)^d}{\alpha_0} \leq \phi_d(\Delta t).$$

To prove (4.6) the reasoning is the same yielding this time,

$$\sum_k \frac{1}{(1 - |S(s_k)|)^2} \leq \frac{4}{c^2 \Delta t^2} \sum_k \frac{1}{\|k\|_2^4} + \frac{(K-1)^d}{\alpha_0^2} \leq \bar{C}(d, \lambda)^2 \Delta t^{-2}, \quad (\text{A.9})$$

cf. Lemma A.2. To prove (4.7), we also follow the same strategy, giving

$$K^{-d} \sum_k \frac{(1, v_k^h)}{1 - |S(s_k)|} \leq \frac{K^{-d}}{c\Delta t} \sum_k \frac{(1, v_k^h)}{\|k\|_2^2} + \frac{K^{-d}}{\alpha_0} \sum_k (1, v_k^h). \quad (\text{A.10})$$

By the Cauchy-Schwarz inequality, Lemma A.1 and Lemma A.2 we have

$$\begin{aligned} \sum_k \frac{(1, v_k^h)}{\|k\|_2^2} &\leq \left( \sum_k (1, v_k^h)^2 \right)^{1/2} \left( \sum_k \frac{1}{\|k\|_2^4} \right)^{1/2} \leq 2^d \pi^{-d} K^d \left( \sum_{n \text{ odd}} \frac{1}{n^2} \right)^{d/2} \sqrt{\tilde{c}(d)} \\ &< 2^{-d/2} K^d \sqrt{\tilde{c}(d)}, \end{aligned}$$

where we have used the fact that the sum of the inverse squares of all odd numbers is  $\pi^2/8$ . For the second summation in (A.10) we get owing to Lemma A.1,

$$\sum_{\mathbf{k}} (\mathbf{1}, \mathbf{v}_{\mathbf{k}}^h) \leq 2^d \pi^{-d} K^d \left( \sum_{n \text{ odd}} \frac{1}{n} \right)^d \leq \pi^{-d} K^d (\log(K-1) + 2)^d,$$

where we used the fact that the harmonic series over the first  $K-1$  odd numbers is bounded by  $\log(K-1)/2 + 1$ . The two bounds for the summations and (A.10) give

$$K^{-d} \sum_{\mathbf{k}} \frac{(\mathbf{1}, \mathbf{v}_{\mathbf{k}}^h)}{1 - |S(s_{\mathbf{k}})|} \leq 2^{-d/2} c^{-1} \sqrt{\tilde{c}} \Delta t^{-1} + \pi^{-d} (\log(K-1) + 2)^d \leq \tilde{C}(d, \lambda) \Delta t^{-1},$$

which is (4.7) and concludes the proof.  $\square$