

1991

B.P. Sommeijer

Parallel-iterated Runge-Kutta methods for stiff ordinary differential equations

Department of Numerical Mathematics Report NM-R9117 October

CWI, nationaal instituut voor onderzoek op het gebied van wiskunde en informatica

CWI is the research institute of the Stichting Mathematisch Centrum, which was founded on February 11, 1946, as a non-profit institution aiming at the promotion of mathematics, computer science, and their applications. It is sponsored by the Dutch Government through the Netherlands organization for scientific research (NWO).

Parallel-Iterated Runge-Kutta Methods for Stiff Ordinary Differential Equations

B.P. Sommeijer

CWI

P.O. Box 4079, 1009 AB Amsterdam, The Netherlands

Abstract: For the numerical integration of a stiff ordinary differential equation, fully implicit Runge-Kutta methods offer nice properties, like a high classical order and high stage order as well as an excellent stability behaviour. However, such methods need the solution of a set of highly coupled equations for the stage values and this is a considerable computational task. This paper discusses an iteration scheme to tackle this problem. By means of a suitable choice of the iteration parameters, the implicit relations for the stage values, as they occur in each iteration, can be uncoupled so that they can be solved in parallel. The resulting scheme can be cast into the class of Diagonally Implicit Runge-Kutta (DIRK) methods and, similar to these methods, requires only one LU-factorization per step (per processor). The stability as well as the computational efficiency of the process strongly depends on the particular choice of the iteration parameters and on the number of iterations performed. We discuss several choices to obtain good stability and fast convergence. Based on these approaches, we wrote two codes possessing local error control and stepsize variation. We have implemented both codes on an ALLIANT/FX4 machine (four parallel vector processors and shared memory) and measured their speedup factors for a number of test problems. Furthermore, the performance of these codes is compared with the performance of the best stiff ODE codes for sequential computers, like SIMPLE, LSODE and RADAU5.

1980 *Mathematics Subject Classification*: 65M10, 65M20

1982 *CR Categories*: G.1.7

Key Words and Phrases: Parallelism, stiffness, Diagonally Implicit Runge-Kutta methods, stability, convergence

1. INTRODUCTION

Due to the never-ending demand for more speed in scientific computation, the available computerpower of new architectures has tremendously increased during the last decades. This is mainly obtained by new hardware design and by a prodigious progress in micro-electronics. However, this *hardware advancement* is not sufficient to meet the requirements as they occur in large-scale problems. The main problem in effectively exploiting this huge potential of computerpower is the fact that there is very *little software* available for these machines. In order to be efficient, this software should be based on algorithms that are well tuned to the new architectures.

Since many numerical algorithms were designed for the traditional sequential computers, the existing methods are not necessarily the best. This is particularly true in the field of numerical methods for ordinary differential equations. Therefore, it is highly desirable to (re)consider these algorithms and, eventually, replace them with more suitable candidates.

Herewith, we arrive at the major aim of this paper: *the construction of new algorithms*, specifically designed for a wide class of new architectures, thus making an attempt to decrease the arrears of software with respect to hardware.

In this paper we will concentrate on numerical methods for the initial-value problem (IVP) for the ordinary differential equation (ODE), written in the autonomous form

$$(1.1) \quad y'(t) = f(y(t)), \quad 0 \leq t \leq T, \quad y \in \mathbb{R}^N, \quad f: \mathbb{R}^N \rightarrow \mathbb{R}^N.$$

In particular, we shall discuss the construction of algorithms for (1.1) that are suitable in a *parallel* environment. Although parallel computers are available now for quite a few years, it is remarkable that this area received only marginal attention and in fact is still in its infancy. A possible explanation may be that the integration of an IVP by a step-by-step process is sequentially in nature and thus offers little scope to exploit parallelism.

Nevertheless, there are some avenues: at first, there is the rather obvious way to distribute the various components of the system of ODEs amongst the available processors. This is especially effective in *explicit* methods, since they frequently need the evaluation of the right-hand side function f for a given vector y , so that the components of f can be evaluated independently of one another. Following the terminology of Gear [11], this is called *parallelism across the problem*. A more interesting approach, called *parallelism across the method*, is to employ the parallelism inherently available within the method. Concurrent evaluations of the entire function f for various values of its argument, and the simultaneous solution of various (nonlinear) systems of equations are examples of parallelism across the method. Remark that this form of parallelism is also effective in case of a scalar ODE (i.e., $N=1$ in (1.1)), whereas parallelism across the problem aims at large N -values. Also notice that both approaches can be combined because they are more or less 'orthogonal'. Still another approach, which could be termed *parallelism across the time*, is followed by Bellen et al. [2]. Contrary to the step-by-step idea, they perform a number of steps simultaneously, thus calculating numerical approximations in many points on the t -axis in parallel. In fact, these methods belong to the class of *waveform relaxation* methods. These methods show a significant speedup provided that the number of steps is (very) large. In the present paper we will confine ourselves to parallelism across the method.

Unfortunately, many existing algorithms that perform well on a sequential computer can take hardly profit from a parallel configuration. This feature necessitates us to construct *new* methods, specifically designed for parallel execution. In doing so, it was in many cases unavoidable to introduce some redundancy in the total volume of computational arithmetic. As a consequence, it is overambitious to expect a speedup (compared with a good sequential solver) in the solution time with a factor s , if s processors are available.

In many of the methods considered in this paper, a *small* number (typically in the range from 2 to 6) of concurrent subtasks of *considerable* computational complexity can be distinguished. Consequently, (i) these methods are aiming at so-called 'coarse-grain' parallelism and (ii) communication and synchronization overhead will be small compared with CPU time.

2. RUNGE-KUTTA METHODS

The general Runge-Kutta (RK) method to proceed the numerical solution of (1.1) from t_n over a step h is given by

$$(2.1a) \quad y_{n+1} = y_n + h \sum_{i=1}^s b_i f(Y_i),$$

$$(2.1b) \quad Y_i = y_n + h \sum_{j=1}^s a_{ij} f(Y_j), \quad i = 1, \dots, s.$$

Here, $y_n \approx y(t_n)$, a_{ij} , b_i are the coefficients defining the RK method and s is called the number of stages. The quantities Y_i , the stage-values, can be considered as intermediate approximations to the solution y . An RK method is said to be explicit iff $a_{ij}=0$, $j \geq i$. Otherwise, it is called an implicit RK (IRK) method. For the algorithms described in this paper, our starting point will always be an IRK method.

A nice feature of IRK methods is that a high order of accuracy can be combined with excellent stability properties [6]. Well-known examples of such IRKs are the Gauss-Legendre methods (order $2s$ and A-stable) and the Radau IIA methods (order $2s-1$ and L-stable). A serious disadvantage however, is the high cost of solving the algebraic equations defining the stage-values Y_i . Since the Y_i are coupled in general, this is a system of dimension $s \cdot N$, thus involving $O((s \cdot N)^3)$ arithmetic operations. This compares unfavourably with ODE solvers based on linear multistep (LM) methods, where a system of dimension N has to be solved in each step. This is the main reason that IRK methods have not received great popularity to serve as the basis for efficient, production oriented software. In the literature, several remedies have been proposed to reduce the amount of linear algebra per step, like Diagonally Implicit RK (DIRK) methods [24, 8, 1, 7] and Singly Implicit RK (SIRK) methods [3, 5]. However, both approaches have their own disadvantages and did not succeed in completely superseding the LM-type of methods. Another possibility to realize the excellent prospects that IRK methods offer, is the use of parallel processors. In [17], [18] and [19], we analyzed several parallel methods. In the next sections we will summarize the main characteristics.

Motivated by our starting point that parallelism across the method should also be effective for scalar ODEs, we will assume throughout that (1.1) is a *scalar* equation. This has the notational advantage that we can avoid tensor products in our formulation. However, the extension to systems of ODEs, and therefore to nonautonomous equations, is straightforward.

In describing the parallel methods, it will be convenient to use a compact notation for the RK method (2.1). Introducing $A=(a_{ij})$, $b=(b_i)$, $Y=(Y_i)$ and $e=(1, \dots, 1)^T$, all of dimension s , a succinct notation of the RK method reads

$$(2.2a) \quad y_{n+1} = y_n + h \, b^T f(Y),$$

$$(2.2b) \quad Y = y_n \, e + h \, A f(Y),$$

where $f(v):=(f(v_j))$, for a given vector $v=(v_j)$.

3. DIAGONAL ITERATION

The main problem in the application of an IRK is the solution of (2.2b) for the stage vector \mathbf{Y} ; once this vector has been obtained, (2.2a) is straightforward. A direct treatment to solve (2.2b) (i.e., applying some form of modified Newton iteration) offers little scope to exploit parallelism, except for the linear algebra part; this aspect is not discussed here, since it is 'orthogonal' to the subject of this paper, i.e., the parallel calculation of the stage vector \mathbf{Y} . To that purpose, we introduce the iteration process

$$(3.1a) \quad \mathbf{Y}^{(j)} - h \mathbf{D} f(\mathbf{Y}^{(j)}) = \mathbf{y}_n \mathbf{e} + h [\mathbf{A} - \mathbf{D}] f(\mathbf{Y}^{(j-1)}), \quad j=1, \dots, m.$$

Here, \mathbf{D} is a *diagonal* matrix. This is crucial, since now, given an iterate $\mathbf{Y}^{(j-1)}$, each individual component $Y_i^{(j)}$ of the unknown iterate $\mathbf{Y}^{(j)}$ has to be solved from an implicit relation of the form

$$(3.1b) \quad Y_i^{(j)} - h d_i f(Y_i^{(j)}) - \Sigma_i = 0, \quad i=1, \dots, s,$$

where Σ_i is the i th component of the right-hand side vector in (3.1a) and d_i is the i th diagonal entry of the matrix \mathbf{D} . Clearly, all Σ_i depend on $\mathbf{Y}^{(j-1)}$, but can be computed straightforwardly (even in parallel). The bulk of the computational effort involves the solution of the s equations for the components $Y_i^{(j)}$, $i=1, \dots, s$. However, given the Σ_i , the equations (3.1b) are *uncoupled* and can be solved in parallel. Hence, assuming that we have s processors available, each iteration in (3.1a) requires *effectively* the solution of only one implicit relation of the form (3.1b). This is especially advantageous in case of (large) systems of ODEs, because then each iteration in (3.1a) requires effectively the solution of a system of dimension N , the ODE dimension. As a consequence, the total iteration process has the effect that the solution of one system of dimension $s \cdot N$ has been transformed into the solution of a sequence of m systems, all of dimension N . Moreover, since \mathbf{D} is the same in all iterations, the (parallel) LU-decompositions of the matrices $\mathbf{I} - h d_i \partial f / \partial \mathbf{y}$ can be restricted to the first iteration. Summing up, the total computational complexity of the iteration process is $O(N^3 + mN^2)$, whereas a direct treatment requires $O(s^3N^3 + MN^2)$, with M the number of (modified) Newton iterations required. Since typical s -values range from 2 to 6 and because the required number of iterations m is quite modest (as we shall see in the Sections 3.1 and 3.2), we now arrive at a manageable level of arithmetic. Notice that this approach is quite similar to that of a DIRK method, where also only one LU-decomposition of a matrix of dimension N is required per step. In [24, 8, 1], A-stable DIRKs are analyzed of order p with $p-1$ implicit stages ($2 \leq p \leq 4$). Cooper and Sayfy [7] have constructed A-stable DIRKs with 5 implicit stages. They present a method of order 5 and could increase the order to 6 by adding one explicit stage. We are not aware of A-stable DIRKs of higher order. However, the parallel approach allows for A-stable methods of as high an order as 10 (excluding 9) (cf. Section 3.1) or even arbitrary high order (cf. Section 3.2).

A further advantage of the parallel methods is that the stage order [9] can be made higher than that of a classical DIRK method. We postpone the discussion of this aspect until Section 3.2 and first finish the discussion of the iteration scheme (3.1a).

To start the iteration (3.1a), we need the initial approximation $\mathbf{Y}^{(0)}$. One of the possibilities to choose this vector is given by

$$(3.1c) \quad \mathbf{Y}^{(0)} - h \mathbf{B} \mathbf{f}(\mathbf{Y}^{(0)}) = y_n \mathbf{e} + h \mathbf{C} \mathbf{f}(y_n \mathbf{e}).$$

Here, the matrix \mathbf{B} will be chosen either zero or of diagonal form in order to exploit parallelism (in the same way as described for (3.1a)); \mathbf{C} is an arbitrary full matrix. Particular choices of these matrices will be discussed in the Sections 3.1 and 3.2. In the sequel, the initial approximation $\mathbf{Y}^{(0)}$ will be referred to as the *predictor*.

If m iterations have been performed with (3.1a), then the new approximation at t_{n+1} is defined by (cf. (2.2a))

$$(3.2a) \quad y_{n+1} := y_n + h \mathbf{b}^T \mathbf{f}(\mathbf{Y}^{(m)}).$$

Once an underlying IRK has been selected (henceforth called the *corrector*), the freedom left in the iteration process (3.1) consists of the matrices \mathbf{B} , \mathbf{C} and \mathbf{D} , and the number of iterations m . With respect to the matrix \mathbf{D} we have considered several possibilities: first of all, there is the simplest choice, which sets \mathbf{D} equal to the zero-matrix. In this case we obtain an *explicit* iteration process and, consequently, the resulting scheme is only suitable for *nonstiff* equations. This approach has received relatively much attention in the literature (see e.g. [26, 23, 20, 4, 21, 16]). Choosing the 'trivial' predictor $\mathbf{Y}^{(0)} = y_n \mathbf{e}$, the order behaviour of the resulting algorithm can be formulated as (see also [20, 21, 22])

Theorem 3.1. The method $\{(3.1a) \text{ with } \mathbf{D}=\mathbf{O}, (3.1c) \text{ with } \mathbf{B}=\mathbf{C}=\mathbf{O}, (3.2a)\}$ is of order $\min\{p^*, m+1\}$, where p^* is the order of the corrector (2.2) []

Notice that this method is itself an explicit RK methods with $s-m+1$ stages. However, on a parallel machine, the *effective* number of stages equals only $m+1$ (provided that s processors are available). This means that if the number of iterations $m \leq p^*-1$, then we have obtained an explicit RK method where the number of effective stages equals the order. This is an optimal result [20] and compares favourably with the situation for classical (uniprocessor) explicit RK methods, where the number of stages increases faster than linearly if we want a high order. Furthermore, we observe that the number s of required processors is minimal with respect to the order, if the generating RK method is of Gauss-Legendre type, since these methods have the highest possible order with respect to the number of stages. We do not continue the discussion of the case $\mathbf{D}=\mathbf{O}$ since this paper aims at *stiff* problems, leading us to *implicit* methods, i.e., to the case $\mathbf{D} \neq \mathbf{O}$.

Before specifying particular choices of \mathbf{D} , we first want to discuss an aspect of the corrector which is relevant with respect to stiffness. In integrating stiff ODEs, a favourable property of the method is that it is 'stiffly accurate'. This notion has been introduced by Prothero and Robinson [27] and means that the RK method satisfies $\mathbf{b}^T = \mathbf{e}_s^T \mathbf{A}$, with \mathbf{e}_s the s th unitvector. Hence, \mathbf{b}^T equals the last row of \mathbf{A} ,

or equivalently, the last component of the stage vector \mathbf{Y} is an approximation to the solution at the new steppoint t_{n+1} . Therefore, in case of a stiffly accurate corrector, (3.2a) will be replaced by

$$(3.2b) \quad y_{n+1} := \mathbf{e}_s^T \mathbf{Y}^{(m)}.$$

Another aspect worth mentioning is, that - for nonstiffly accurate correctors - the final evaluation $f(\mathbf{Y}^{(m)})$ in (3.2a) has a bad influence on the stability of the iterated scheme [28]. This can be avoided by the following modification (see [12]): suppose that the corrector would have been solved, i.e., \mathbf{Y} satisfies (2.2b). Then (assuming that A is nonsingular), we have

$$(3.3) \quad h f(\mathbf{Y}) = A^{-1} [\mathbf{Y} - y_n \mathbf{e}].$$

Replacing \mathbf{Y} in this relation by $\mathbf{Y}^{(m)}$ (that is, assuming that (3.3) is reasonably satisfied by $\mathbf{Y}^{(m)}$) and substitution into (3.2a) leads to

$$(3.2c) \quad y_{n+1} := y_n + \mathbf{b}^T A^{-1} [\mathbf{Y}^{(m)} - y_n \mathbf{e}].$$

In case of a nonstiffly accurate corrector, the use of (3.2c) instead of (3.2a) has two consequences for the resulting method: the stability is improved by using (3.2c), since we avoid the final evaluation $f(\mathbf{Y}^{(m)})$; on the other hand, (3.2a) is more accurate (see also the remark following Theorem 3.2). For a stiffly accurate corrector, however, (3.2a) and (3.2c) are equivalent.

Now, we return to the discussion of the matrix D ; we distinguish two cases:

- (i) D is such that after a *prescribed* number of iterations the resulting method has good stability properties. This option has been followed in [17] and will be outlined in Section 3.1. In this approach the order of the resulting method equals the order of the corrector and the number of iterations m is minimal to reach this order.
- (ii) Another option is to *solve* the corrector and to choose D in such a way that we obtain fast convergence in the iteration process (3.1a). This strategy has been followed in [18] and [19] and will be the subject of Section 3.2.

In the following sections these cases will be briefly discussed; henceforth, the above Parallel Diagonally-Iterated RK methods will be denoted by PDIRK methods.

3.1. Diagonal iteration with a prescribed number of iterations.

Here, we consider methods for which the number of iterations m will be fixed. As we shall see, this number is dictated by the orders of the corrector and of the predictor. To clarify this strategy, we quote a theorem from [17]:

Theorem 3.2. Let p^* be the order of the underlying corrector (2.2). Then the order p of the resulting PDIRK method $\{(3.1), (3.2a), (3.2b)\}$ is given by

$$\begin{aligned}
&\min \{p^*, m+r\} && \text{for all matrices } B, C \text{ and } D, \\
&\min \{p^*, m+1+r\} && \text{if } (C+B)e = Ae, \\
&\min \{p^*, m+2+r\} && \text{if, in addition, } BAe = A^2e,
\end{aligned}$$

where r takes the value 1 if y_{n+1} is defined by (3.2a) (i.e., the nonstiffly accurate case) and $r=0$ if y_{n+1} is defined by (3.2b) (the stiffly accurate case).

Furthermore, if the corrector is stiffly accurate, then the corresponding PDIRK method has the same property. \square

Remark: For the nonstiffly accurate case, we observe that if we change to the modification (3.2c), then r should be set to 0 in Theorem 3.2. Since in [17] and [18], the nonstiffly accurate methods were analyzed on the basis of (3.2a), we will confine ourselves in this overview to this choice.

Based on this theorem, we adopted in [17] the strategy to choose m such that $p=p^*$. This means that we stop iterating as soon as the order has reached the order of the corrector, since a continuation of the iteration process would not increase the order of the PDIRK method. Furthermore, in [17], we only considered correctors of Gauss-Legendre type and of Radau IIA type.

With respect to the choice of the predictor, it turned out that employing the C -matrix did not yield particular advantage; so, here we only present results for $C=O$. For the matrix B we choose either $B=O$ or $B=D$. Although B and D may be different diagonal matrices, the choice $B=D$ has the computational advantage that the LU-decompositions of $I-d_i h \partial f / \partial y$, which are needed during the iteration (3.1a), can also be used in solving (3.1c) for $Y^{(0)}$.

The diagonal matrix D is chosen such that the resulting PDIRK method has optimal stability characteristics. Here, we distinguish two approaches: matrices D with constant and with varying diagonal entries. These variants will be discussed in the following subsections, respectively.

3.1.1. D-matrices of the form $d \cdot I$. In this relatively simple case we could perform a rather thorough stability analysis, using the so-called 'E-polynomials' (see e.g. [6]). In this connection we also mention the work of Wolfbrandt [29], who investigated similar stability functions. A few classes of unconditionally stable methods are listed in Table 3.1. The values of d can be found in [17].

Table 3.1. Unconditionally stable PDIRK methods with $D=d \cdot I$

corrector	matrices B and D	attainable order p	# effective stages	stability
Gauss	$B = O, D = d \cdot I$	$p \leq 4$	$p-1$	A-stable
Gauss	$B = D = d \cdot I$	$p \leq 6, p = 8$	p	L-stable
Radau IIA	$B = O, D = d \cdot I$	$p \leq 6, p = 8$	p	L-stable
Radau IIA	$B = D = d \cdot I$	$p \leq 8, p = 10$	$p+1$	L-stable

We recall that the Gauss and Radau IIA methods are good choices to serve as a corrector, since these IRKs have a high order with respect to the number of processors required (i.e., these methods need a minimal number of stages). It is however interesting to remark that any RK method can be used as a corrector, even an *explicit* one, although in that case we have the unconventional situation that an explicit corrector is iterated by means of an implicit iteration process. For example, the PDIRK scheme resulting from iterating a p th-order explicit RK method ($p \leq 6$ or $p=8$) using exactly p iterations and $B=D=d \cdot I$, can be made L-stable by choosing the appropriate d -value. However, the number of processors equals the number of stages of the explicit RK method and thus is at least p .

3.1.2. Nonconstant D-matrices. If we omit the restriction of constant elements in D , then we can save one iteration and still obtain the same order as in the previous subsection, simply by setting $B=D=\text{diag}(Ae)$ (cf. Theorem 3.2). The stability function for this case is rather complicated, so that the stability region of the methods had to be determined numerically. Some of these methods turned out to be only $A(\alpha)$ -stable, however with α close to 90° . In Table 3.2, we collect a few methods with good stability properties.

Table 3.2. PDIRK methods with a nonconstant D-matrix

corrector	attainable order p	# effective stages	stability
Gauss/Radau IIA	$p \leq 5$	$p-1$	strongly A-stable
Gauss/Radau IIA	$p = 6, 7$	$p-1$	$A(\alpha)$ -stable, $\alpha > 83^\circ$
Radau IIA	$p = 3, 5, 7$	p	$L(\alpha)$ -stable, $\alpha > 89^\circ$

3.1.3. Numerical example on the ALLIANT FX/4. Here, we will show the performance of an L-stable PDIRK scheme with $B=D=d \cdot I$ (cf. Table 3.1) when running on a parallel computer. Based on the 4-stage Radau IIA method, we perform 7 iterations to arrive at order 7. Hence, including the calculation of the predictor, this PDIRK scheme requires, effectively, 8 stages per step. This method is L-stable for a range of d -values, from which we selected $d=0.1690246379$. This special d -value has the effect that the degree of the denominator in the (rational) stability function is *two* larger than the degree of the numerator, which causes extra damping at infinity.

We equipped this method with a provisional strategy for error control and stepsize selection. Since the PDIRK approach inherently provides a whole set of *embedded reference solutions* of lower order, a simple way to obtain an estimate for the local truncation error is given by $\|e_s^T Y^{(m)} - e_s^T Y^{(j)}\|$ for some $j < m$. Notice that this estimate does not require additional computations, since $Y^{(j)}$ is anyhow needed to proceed the iteration process. In our code we set $j=m-1$ (recall that $s=4$ and $m=7$). For further details concerning the strategy, we refer to [17].

We implemented this scheme on the ALLIANT FX/4 computer (four parallel processors and shared memory) and applied it to several test problems. The goal of these tests is twofold: (i) we want to investigate to what extent the theoretical parallelization can be realized in practice; in other words, how close we can approach the ideal speedup factor on this four-processor machine and (ii) we want

to compare the performance of the code PDIRK with that of a good sequential solver. To that purpose we selected the reliable code SIMPLE of Nørsett and Thomsen [25] which is based on an A-stable DIRK method. Its (fixed) order is 3, which is rather low. Moreover, it has stage order 1. Since many problems are more efficiently integrated if high-order formulas are available, we also included in our tests the code LSODE of Hindmarsh [14]. This BDF-based code has formulas up to order 5 available, from which only those of first and second order are A-stable. Hence, LSODE is less robust as a *general* stiff solver, but, on the other hand, is generally accepted as a good sequential solver and enjoys considerable usage over a long period.

One of the test problems described in [17] is the set of reaction rate equations:

$$(3.4) \quad \begin{aligned} \frac{dy_1}{dt} &= -0.04 y_1 + 10^4 y_2 y_3, & y_1(0) &= 1, \\ \frac{dy_2}{dt} &= 0.04 y_1 - 10^4 y_2 y_3 - 3 \cdot 10^7 (y_2)^2, & y_2(0) &= 0, & 0 \leq t \leq 10^8, \\ \frac{dy_3}{dt} &= 3 \cdot 10^7 (y_2)^2, & y_3(0) &= 0. \end{aligned}$$

This problem is also used by Nørsett & Thomsen and Hindmarsh to illustrate the codes SIMPLE and LSODE. Initially, the solution changes rapidly and small stepsizes are necessary; gradually, the problem reaches a steady state and an efficient integration requires the stepsize to be increased significantly. In a typical situation, we observed stepsizes in the range $[10^{-3}, 10^6]$. For several values of TOL (the local error bound), the results of the various codes are given in Table 3.3. Here, T_1 and T_4 denote the CPU-time (in seconds) when the program is run on 1 and 4 processors, respectively. The accuracy is measured by means of Δ , which is defined by writing the maximum norm of the global (absolute) error in the end point in the form $10^{-\Delta}$.

Table 3.3. Performance of the codes SIMPLE, LSODE and PDIRK for problem (3.4)

Method	(effective) number of implicit relations/step	TOL	Δ	T_1	T_4
SIMPLE	3	10^{-4}	6.5	0.63	0.85
		10^{-5}	7.8	1.38	$>T_1$
		10^{-6}	9.5	3.67	$>T_1$
LSODE	1	10^{-5}	7.4	0.35	$>T_1$
		10^{-7}	8.6	0.80	$>T_1$
		10^{-9}	10.3	1.71	$>T_1$
PDIRK	8	10^2	8.5	0.51	0.19
		10^0	11.1	1.08	0.37

From this experiment we can conclude:

- (i) concerning the parallelization of the PDIRK code we observe a speedup (defined by T_1/T_4) with a factor ≈ 2.8 . The main reason for not obtaining the ideal factor 4 is that the various processors needed a different number of Newton iterations to solve their 'own' implicit relations. We counted the total number of Newton iterations (over the whole integration interval) for each individual processor and observed that the two extreme values differ by about 20%.
- (ii) the scalar codes SIMPLE and LSODE run faster on one processor than on four. Apparently, the parallelization overhead degrades the performance. Moreover, the dimension of the ODE (3.4) is too small to take any advantage from the vectorization capabilities of the ALLIANT.
- (iii) when compared with PDIRK, we see that SIMPLE needs much more time in the high-accuracy range. This is obviously due to its low order. LSODE is more efficient in this range but, when compared to PDIRK, its CPU-time is approximately four times as larger to obtain 8.5 digits (absolute) precision.
- (iv) finally, we observe that the value for TOL used by PDIRK is several orders of magnitude larger than the value used by either SIMPLE or LSODE to achieve the same global error. This can be explained as follows: Due to its high order, the local truncation error of PDIRK is usually relatively small. Therefore, if crude tolerances are used, the error control mechanism signals that a large stepsize can be used in order to balance the estimated and the requested local error. On the other hand, the Newton process imposes a limitation on the stepsize. In our implementation, the Newton processes to solve for (the components of) $Y^{(0)}$ are given the value y_n as initial iterate. Unfortunately, for large values of h (as suggested by the error estimator) this initial iterate is not always inside the contraction domain for the Newton process, resulting in an adequate reduction of the stepsize. As a consequence, this high-order scheme, using a small(er) stepsize, will produce a local error which is much smaller than requested. In conclusion, for this test problem, the restriction on the stepsize imposed by the Newton process is more stringent than that imposed by the local error control, unless very small values for TOL are used. We have also integrated some *linear* ODEs (for which the convergence problems are not relevant, of course) and observed a relation between TOL and the global error similar to that of SIMPLE and LSODE.

3.2. Diagonal iteration until convergence.

PDIRK methods with a fixed number of iterations, as considered in the previous subsection, are in fact special DIRK methods. It is well known [9] that DIRK methods possess a low so-called *stage order* (viz., 1) which, in general, drastically reduces the accuracy. As a matter of fact, in many stiff problems the actually observed order equals the stage order (or, sometimes the stage order+1). As a consequence of this so-called *order-reduction* phenomenon, the relevance of methods with a high algebraic (i.e., classical) order and a low stage order is questionable. Therefore, apart from the 'fixed-m-strategy' we also considered the approach where the corrector is iterated *until convergence*. This implies that we can rely on all the characteristics of the corrector, like stability and accuracy behaviour and, in particular, the stage order. For example, s-stage IRK methods of Gauss and Radau type both have stage order s . In addition, they have a very high algebraic order (superconvergence) but, as observed above, this property seems to be of minor importance in many stiff problems.

Therefore, we also considered (A-stable) Newton-Cotes and Lagrange type IRKs (cf. [18]); in these (collocation) methods the superconvergence is exchanged for an increase by one of the stage order. This is obtained by adding one explicit stage to the s implicit stages. The time needed for this extra explicit stage is quite negligible compared with the time involved in solving the implicit stages. Thus, we arrive at correctors with algebraic order = stage order = $s+1$, which are suitable for parallel iteration on an s -processor machine.

Having decided to *solve* the corrector, we can now consider (3.1a) as an *iteration* process, where 'iteration' has the classical meaning. This leads us automatically to a criterion for choosing the matrix D : this matrix should be such that we have *fast convergence* in (3.1a).

In [18] it was shown that the iteration error $Y - Y(j)$, in first approximation, satisfies the recursion

$$(3.5a) \quad Y - Y(j) = Z(z) [Y - Y(j-1)], \quad j=1, \dots, m, \quad z := h \lambda,$$

where the iteration matrix Z is defined by

$$(3.5b) \quad Z(z) := zD[I - zD]^{-1}[D^{-1}A - I].$$

Here, λ denotes an approximation to the derivative $\partial f/\partial y$ and should be understood to run through the spectrum of the Jacobian matrix in case of systems of ODEs. The convergence behaviour of (3.1a) is completely determined by the iteration matrix Z and we have the matrix D at our disposal to obtain fast convergence.

The main difficulty in choosing D is that Z depends on z , i.e., on the problem. Therefore, we cannot expect to find a uniformly 'best' D -matrix. Since we are aiming at the integration of stiff equations, we considered the influence of Z on the eigenvectors of $\partial f/\partial y$ corresponding to eigenvalues of large modulus. For $|z| \rightarrow \infty$, Z behaves as $I - D^{-1}A$. Thus a strong damping of these eigenvectors leads us to the minimization of the spectral radius of $I - D^{-1}A$. Observe, that the 'nonstiff' eigenvectors (corresponding to small values of $|z|$) are already damped since Z behaves as $z[A - D]$ for $|z| \rightarrow 0$. With this approach we obtained fast convergence on a broad collection of test examples (cf. [18]). However, we do not claim that this choice of D is the best possible. For example, a more sophisticated strategy might be the minimization of (some norm of) $Z(z)$ over the whole, or the 'stiff part' of the left halfplane.

Another possibility could be to minimize the principal stiff error constants in the resulting PDIRK method; this option has been studied in [19]. Several other options to choose D are discussed in [18]. Many of these have been used in numerical tests, but it turned out that the behaviour of the strategy based on the minimization of the spectral radius ρ of $I - D^{-1}A$ could not be improved.

Based on this approach, methods have been constructed for $s=2,3,4$ (cf. [18]). Only for $s=2$ it is possible to determine D analytically such that $\rho(I - D^{-1}A) = 0$. For the larger values of s , the D -matrices had to be calculated numerically. The ρ -values found, increase with s and are (for the several correctors) in the range (0.004, 0.01) if $s=3$ and in the range (0.02, 0.1) for $s=4$.

In [18], we also made a mutual comparison of several *stiffly accurate* correctors. It turned out that, in general, the Radau IIA and the Lagrange based methods are superior to the Lobatto IIIA and

Newton-Cotes based methods. This is probably due to the fact that the first methods have damping at infinity, whereas the latter type of methods are only weakly stable at infinity. Furthermore, we also considered the *nonstiffly accurate* method based on the Gauss corrector. This method showed poor stability behaviour, but as mentioned before, this can be improved upon if we use (3.2c) instead of (3.2a). Evidently, the final iteration error $\mathbf{Y} - \mathbf{Y}^{(m)}$ depends on the initial error $\mathbf{Y} - \mathbf{Y}^{(0)}$ (see (3.5a)). In [18], only the 'trivial' predictor $\mathbf{Y}^{(0)} = \mathbf{y}_n \mathbf{e}$ (i.e., $\mathbf{B} = \mathbf{C} = \mathbf{O}$ in (3.1c)) has been used. In [19], also the implicit variant is considered ($\mathbf{B} = \mathbf{D}$, $\mathbf{C} = \mathbf{O}$), as well as predictors that use information from the previous step. This is a natural way to increase the accuracy of the predictor, since all methods are based on the collocation principle. This implies that the stage vector $\mathbf{Y}^{(m)}$ calculated in the preceding step defines a collocation polynomial which can be extrapolated to the present step. Needless to say that, in general, the increased accuracy of $\mathbf{Y}^{(0)}$ results in fewer iterations.

Apart from the convergence behaviour we also studied the stability of the iterated methods for several, but fixed values of m . It can be shown that the stability functions of the methods based on the Lobatto IIIA and Newton-Cotes correctors (using $\mathbf{Y}^{(0)} = \mathbf{y}_n \mathbf{e}$) are only A-acceptable in the limit, i.e., for $m \rightarrow \infty$ (cf. [18, Table 4.1]). For the Radau IIA and Lagrange based methods the stability function of the PDIRK method is already A-acceptable for modest values of m . These m -values for various choices of $\mathbf{Y}^{(0)}$ can be found in [19, Table 3.2a]. However, for many problems the requirement of A-stability is unnecessarily strong and can be weakened to $A(\alpha)$ -stability with α sufficiently large. It turns out that the angles α corresponding to the Radau IIA and Lagrange correctors are close to 90° after only a few iterations, especially if $\mathbf{Y}^{(0)}$ is defined implicitly ($\mathbf{B} \neq \mathbf{O}$); these α -values can be found in [19, Table 3.2b].

3.2.1. Numerical example on the ALLIANT FX/4. Similar to the 'fixed- m -strategy' discussed in Section 3.1, we have implemented a PDIRK method based on the 'minimal-spectral-radius-strategy'. For the corrector, we selected the 4-stage Radau IIA method; the predictor $\mathbf{Y}^{(0)}$ is obtained from extrapolation of the collocation polynomial calculated in the preceding step. This pair is recommended in [19] as the most efficient combination.

For the error control, we calculate a reference solution of the form

$$(3.6) \quad y_{\text{ref}} = \alpha y_n + \beta_0 h f(y_n) + \sum_{i=1}^4 \beta_i Y_i^{(m)},$$

where the $Y_i^{(m)}$ are the final iterates, and the coefficients in (3.6) are determined to make this reference solution fourth-order accurate. This requirement leaves one coefficient free, say β_0 , which is set to 0.1. Following an idea of Shampine, the local truncation error is estimated by

$$(3.7) \quad \text{local error} \approx \| (\mathbf{I} - d_4 h \partial f / \partial \mathbf{y})^{-1} (y_{\text{ref}} - y_{n+1}) \|,$$

where d_4 is the last entry of the diagonal matrix \mathbf{D} . The premultiplication with $(\mathbf{I} - d_4 h \partial f / \partial \mathbf{y})^{-1}$ in (3.7) serves to obtain a bounded estimate if $h\lambda \rightarrow \infty$ for problems of the form $\mathbf{y}' = \lambda \mathbf{y}$ (see also [12, p.134]).

The resulting code is termed PSODE. In contrast to the code PDIRK, where we used a fixed number of iterations, PSODE is equipped with a strategy to terminate the iteration (3.1a). The stopping criterion for this iteration is related to TOL and a test on the rate of convergence is performed in each iteration. If this test predicts that it is unlikely that convergence will be obtained within ITER_{\max} (in PSODE set to 10) iterations, then the process is interrupted and restarted with a smaller stepsize.

There is however another difference between the two codes, which is of greater impact. In PDIRK, each iterate $\mathbf{Y}^{(j)}$ in (3.1a) is solved up to machine precision by a modified Newton process (similar to the approach followed in conventional DIRK methods). Especially for small j -values, this is a waste of effort, since $\|\mathbf{Y} - \mathbf{Y}^{(j)}\|$ will be relatively large at the start of the iteration process. Moreover (and this is the essential difference between PDIRK methods and conventional DIRKs), $\mathbf{Y}^{(j)}$ is no longer needed once we have calculated $\mathbf{Y}^{(j+1)}$, because both are approximations to the exact solution at the same points.

Therefore, each implicit relation in PSODE is 'solved' by just *one* (modified) Newton iteration. As a result, the number of iterations in (3.1a) to solve (2.2b) will increase, but it turned out that the overall process is more efficient. An additional advantage of this new approach is that we now obtain a perfect load-balancing, since all processors perform exactly the same number of Newton iterations. Hence, a degradation of the performance as observed for the code PDIRK is avoided.

We have implemented the code PSODE on the ALLIANT FX/4 and applied it to a number of test problems. Again, the aim of these test is to measure the speedup of the code on this parallel machine and, additionally, to compare its performance with a good sequential solver. Since for many problems, LSODE is an efficient stiff solver on sequential machines, this code is again an obvious reference method. Furthermore, we selected the recent (sequential) code RADAU5 of Hairer & Wanner [12]. This choice is motivated by the observation that it solves a Radau IIA method (viz. the 3-point 5th-order one); this starting point is quite similar to that of PSODE, although the approach to obtain the Radau-solution is completely different. Since the DIRK-based code SIMPLE is of a different nature and also because of its inefficient behaviour in the high-accuracy range, we decided to cancel this code as a reference method.

In comparing the parallel code PSODE with the two sequential codes, we do *not* take into account effects originating from a possible 'parallelization over the loops'. By this we mean that a long loop is cut into s smaller parts which are then assigned to the s processors. In the Introduction, this effect is termed 'parallelism across the problem' and can in fact be used by any ODE solver. Here we merely want to test *intrinsic* parallelism (called 'parallelism across the method'). In order to exclude the effects of 'parallelism across the problem', LSODE and RADAU5 are run on a single processor. In fact, the amount of intrinsic parallelism offered by LSODE and RADAU5 is very modest (see also the remark at the end of this section).

Of course, if one is interested in 'parallelism across the problem', then the sequential codes could be implemented on an s -processor machine. However, in that case a fair comparison would require assigning 4s processors to PSODE, since in each of the 4 concurrent subtasks of PSODE, the 'parallelism across the problem' can equally well be exploited.

Summarizing, we may say that PSODE needs 4 times the number of processors given to a sequential code, simply because it possesses a 4-fold amount of intrinsic parallelism. The large number of processors utilized by PSODE reflects the current tendency in parallel computing, since modern architectures – and certainly those entering the market in the coming years – have an ‘almost unlimited’ number of processors (*massive parallelism*).

Another aspect which is of utmost importance for the performance of a stiff code, is the amount of linear algebra per step, which in turn strongly depends on the dimension of the ODE. Prior to the specification of our test problem, we will briefly discuss the characteristics of the various codes with respect to this aspect:

A common feature of the three codes is that they need from time to time an LU-decomposition of the matrix involved in their respective iteration processes to solve the nonlinear relations. Since the factorization of a general N -dimensional matrix requires approximately $N^3/3$ arithmetic operations, this will dominate the total costs of the integration for *large-scale problems*. Here we may think of complicated problems from circuit analysis or semi-discretized (higher-dimensional) partial differential equations. In such applications, systems of ODEs with several thousands of equations are quite usual. In this connection we remark that both LSODE and PSODE deal with matrices of dimension N . Hence, it is to be expected that their mutual comparison is only marginally influenced if N increases and all other aspects are left unchanged.

Matters are different for the code RADAU5, since it has to deal with matrices of dimension $3N$. By exploiting the special structures in these matrices, Hairer and Wanner are able to reduce the total work of the LU-decomposition to $5N^3/3$ operations [12], thus gaining a factor 5 compared with a direct treatment, which would have required $(3N)^3/3$ operations. However, this number $5N^3/3$ compares unfavourably with the number $N^3/3$ (associated with LSODE and PSODE), and causes a serious drawback for RADAU5 when applied to large-scale problems.

To get some insight in the performances of the codes, we have applied them to a small test problem originating from circuit analysis. It was first described by Horneber [15] and extensively discussed in [13, p.112] and [10]. This (stiff) system describes a ring modulator, which mixes a low frequency and a high frequency signal. The modulated signal is then used as input for an amplifier.

The resulting system of 15 ODEs is defined by

$$\begin{aligned}
 (3.8a) \quad & y_1' = C^{-1} [y_8 - 0.5 y_{10} + 0.5 y_{11} + y_{14} - y_1/R] \\
 & y_2' = C^{-1} [y_9 - 0.5 y_{12} + 0.5 y_{13} + y_{15} - y_2/R] \\
 & y_3' = C_s^{-1} [y_{10} - g(z_1) + g(z_4)] \\
 & y_4' = C_s^{-1} [-y_{11} + g(z_2) - g(z_3)] \\
 & y_5' = C_s^{-1} [y_{12} + g(z_1) - g(z_3)] \\
 & y_6' = C_s^{-1} [-y_{13} - g(z_2) + g(z_4)] \\
 & y_7' = C_p^{-1} [-y_7/R_i + g(z_1) + g(z_2) - g(z_3) - g(z_4)] \\
 & y_8' = -L_h^{-1} y_1 \\
 & y_9' = -L_h^{-1} y_2 \\
 & y_{10}' = L_s^{-1} [0.5 y_1 - y_3 - 17.3 y_{10}] \\
 & y_{11}' = L_s^{-1} [-0.5 y_1 + y_4 - 17.3 y_{11}] \\
 & y_{12}' = L_s^{-1} [0.5 y_2 - y_5 - 17.3 y_{12}] \\
 & y_{13}' = L_s^{-1} [-0.5 y_2 + y_6 - 17.3 y_{13}] \\
 & y_{14}' = L_t^{-1} [-y_1 + e_1(t) - 86.3 y_{14}] \\
 & y_{15}' = L_t^{-1} [-y_2 - 636.3 y_{15}],
 \end{aligned}$$

where

$$\begin{aligned}
 (3.8b) \quad & z_1 := y_3 - y_5 - y_7 - e_2(t), & z_2 := -y_4 + y_6 - y_7 - e_2(t), \\
 & z_3 := y_4 + y_5 + y_7 + e_2(t), & z_4 := -y_3 - y_6 + y_7 + e_2(t),
 \end{aligned}$$

and the function g , which models the characteristics of the diodes, is defined by

$$(3.8c) \quad g(z) := 40.67286402 \cdot 10^{-9} [\exp(17.7493332 \cdot z) - 1].$$

The signals e_1 and e_2 are defined by

$$(3.8d) \quad e_1(t) := 0.5 \sin(2 \cdot 10^3 \pi t), \quad e_2(t) := 2 \sin(2 \cdot 10^4 \pi t).$$

The technical parameters have been given the values $C=16 \cdot 10^{-9}$, $R=25000$, $C_p=10^{-8}$, $R_i=50$, $L_h=4.45$, $L_s=0.0005$ and $L_t=0.002$, resulting in a heavily oscillating solution. Not yet fixed is the value of the capacity C_s . In our test, we give it the value 10^{-9} , which seems technically meaningful. It is reported [13] that small C_s -values cause serious difficulties. In the limit, i.e. on setting $C_s=0$, we end up with a differential-algebraic system. The integration interval in our test is $[0, 10^{-3}]$; the initial values are given by $y_i(0)=0$, $i=1, \dots, 15$. For several values of the local error bound TOL the results obtained by the codes RADAU5, LSODE and PSODE are collected in Table 3.4. Again, T_1 and T_4 denote the CPU time (in seconds) when the program is run on 1 and 4 processors, respectively. Recall, that we restrict the timings for the sequential codes to T_1 . The accuracy is measured by means of Δ_{rel} , which is defined by writing the maximum norm of the global (*relative*) error in the endpoint in

the form $10^{-\Delta_{\text{rel}}}$. Furthermore, Nsteps denotes the number of (successful) integration steps and \bar{m} stands for the average number of (effective) f-evaluations per step.

Table 3.4. Performance of the codes RADAU5, LSODE and PSODE for the circuit problem (3.8)

Method	TOL	Nsteps	\bar{m}	Δ_{rel}	T_1	T_4
RADAU5	10^{-2}	1275	9.0	1.1	33.1	
	10^{-3}	2277	7.6	2.6	48.6	
	10^{-4}	3922	6.7	3.8	72.4	
	10^{-5}	6761	6.1	4.9	110.9	
LSODE	10^{-3}	7054	1.5	1.4	33.6	
	10^{-4}	9772	1.4	2.8	44.1	
	10^{-5}	13266	1.4	2.9	57.7	
	10^{-6}	17887	1.3	3.8	74.7	
	10^{-7}	23310	1.3	4.5	93.1	
	10^{-8}	30253	1.2	4.9	114.3	
PSODE	10^{-2}	1185	7.3	1.4	80.0	21.4
	10^{-3}	1561	7.3	3.1	104.5	27.8
	10^{-4}	2272	7.1	4.1	146.4	39.6
	10^{-5}	3437	6.9	5.2	212.1	57.7

These results give rise to the following conclusions:

- (i) we see that the speedup factor for PSODE (obviously defined by T_1/T_4) is approximately 3.7, which is pretty close to the ‘ideal’ factor 4 on this machine. This factor rapidly converges to 4 if the dimension of the problem increases.
- (ii) furthermore, we observe a remarkable similarity between RADAU5 and PSODE: both codes need approximately 7 f-evaluations per step; moreover, to produce the same accuracy, the required number of steps is of the same order of magnitude (for the more stringent values of TOL, the difference in the number of steps increases, which is probably due to the higher order of PSODE). There is however a striking difference between the two Radau-based codes and LSODE; this code is very cheap per step, but needs much more integration steps to produce the same accuracy. For example, to obtain a relative accuracy of about 5 digits, PSODE needs ≈ 3400 steps, RADAU5 twice as many, whereas for LSODE this number is 9 times as large. Taking into account the computational effort per step of the various codes, the comparison with PSODE yields a double amount of time both for LSODE and RADAU5. Approximately the same ratios are observed in the low-accuracy range (say, $\Delta_{\text{rel}}=3$).

As mentioned before, this example is only a model problem describing a small (part of an) electrical circuit, and is still far away from a real-life application. However, even for this small system of ODEs, the performance of (this provisional version of) PSODE is already superior by a factor 2 to that of the (well-established) codes LSODE and RADAU5.

Summarizing, we can say that

- the PSODE-approach is much more promising to serve as the basis for an efficient, 'all-purpose' stiff solver than the LSODE-approach. This is due to the improved mathematical qualities, viz. the high order in combination with A-stability.
- In comparison with RADAU5, PSODE has the advantage that in large-scale problems, the (dominating) LU-factorizations require a factor 5 less computational effort. In this connection we remark that a few preliminary experiments with a problem of dimension 75 reveal that the overall gain of PSODE is already more than a factor 4.

For really large-scale problems we expect that the speedup factor will be in the range 6 - 8, depending on the required accuracy. This number is composed of the asymptotic factor 5 coming from the algebra part and the remaining factor 1.2 - 1.6 originating from the higher order of PSODE.

Remark: it should be mentioned that RADAU5 offers a possibility to exploit a small amount of intrinsic parallelism. In using two processors, the total number of arithmetic operations to perform the LU-decomposition can be reduced from $5N^3/3$ to $4N^3/3$. We refrained from adapting the code RADAU5 in order to exploit this feature.

4. CONCLUDING REMARKS

In this paper we proposed an iterative approach to solve the stage vector equations occurring in a fully implicit Runge-Kutta method. By a suitable choice of the iteration parameters, the resulting scheme can be cast into the class of A-stable Diagonally Implicit Runge-Kutta (DIRK) methods. However, the new schemes can be given a much higher order than the classical DIRKs available in the literature. The iterated methods have the special feature that many of the implicit relations can be solved in *parallel*, which offers a great computational advantage. Moreover, because of the 'DIRK-nature' of the new schemes, they require only one LU-factorization (of a matrix with the ODE dimension) per step (per processor).

In this paper we discussed two different iteration strategies and, for both strategies, optimal iteration parameters are derived. For both approaches, a variable stepsize code has been implemented on an ALLIANT FX/4 computer. On the basis of two test problems, the speedup factors have been measured; these factors, which depend on the dimension of the ODE, vary between 2.5 and 3.7, which is pretty close to the 'ideal' factor 4 on this machine.

Furthermore, the performance of the codes is compared with that of the best sequential stiff ODE codes: SIMPLE, LSODE and RADAU5. For a relatively simple ODE of small dimension, the parallel code is slightly more efficient than LSODE and much more efficient than SIMPLE. For a more

difficult problem of a larger dimension (viz., 15 ODEs), the parallel code needs considerably less time than the sequential codes.

Finally, it should be remarked that the parallel codes are still in a research phase and need a further tuning of their parameters on the basis of extensive testing. Furthermore, we plan to extend the codes with the facility to treat ODEs of the form $M \cdot y'(t) = f(y(t))$, where M is a matrix which may be singular, resulting in a differential-algebraic system.

REFERENCES

- [1] **Alexander, R.** (1977): *Diagonally implicit Runge-Kutta methods for stiff ODEs*, SIAM J. Numer. Anal. **14**, pp. 1006-1021.
- [2] **Bellen, A. , Vermiglio, R. & Zennaro, M.** (1990): *Parallel ODE-solvers with stepsize control*, J. Comput. Appl. Math. **31**, pp. 277-293.
- [3] **Burrage, K.** (1978): *A special family of Runge-Kutta methods for solving stiff differential equations*, BIT **18**, pp. 22-41.
- [4] **Burrage, K.** (1989): *The error behaviour of a general class of predictor-corrector methods*, CMSR Report, University of Liverpool.
- [5] **Butcher, J.C.** (1979): *A transformed implicit Runge-Kutta method*, J. Assoc. Comput. Mach. **26**, pp. 731-738.
- [6] **Butcher, J.C.** (1987): *The numerical analysis of ordinary differential equations, Runge-Kutta and general linear methods*, Wiley, New York.
- [7] **Cooper, G.J. & Sayfy, A.** (1979): *Semiexplicit A-stable Runge-Kutta methods*, Math. Comp. **33**, pp. 541-556.
- [8] **Crouzeix, M.** (1975): *Sur l'approximation des équations différentielles opérationnelles linéaires par des méthodes de Runge-Kutta*, Ph. D. Thesis, Université de Paris.
- [9] **Dekker, K. & Verwer, J.G.** (1984): *Stability of Runge-Kutta methods for stiff nonlinear differential equations*, North-Holland, Amsterdam.
- [10] **Denk, G. & Rentrop, P.** (1989): *Mathematical models in electric circuit simulation and their numerical treatment*, in: *Numerical Treatment of Differential Equations*, Halle/DDR 1989, Teubner-Texte zur Mathematik.
- [11] **Gear, C.W.** (1988): *Parallel methods for ordinary differential equations*, Calcolo **25**, pp. 1-20.
- [12] **Hairer, E. & Wanner, G.** (1991): *Solving ordinary differential equations, II: Stiff and differential-algebraic problems*, Springer Series in Comp. Math., Vol. **14**, Springer-Verlag, Berlin.
- [13] **Hairer, E., Lubich, C. & Roche, M.** (1989): *The numerical solution of differential-algebraic systems by Runge-Kutta methods*, Lecture Notes in Mathematics, Vol. **1409**, Springer-Verlag, Berlin.
- [14] **Hindmarsh, A.C.** (1980): *LSODE and LSODI, two new initial value ordinary equation solvers*, ACM/SIGNUM Newsletter **15** (4), pp. 10-11.

- [15] **Horneber, E.H.** (1976): *Analyse nichtlinearer RLCÜ-Netzwerke mit Hilfe der gemischten Potentialfunction mit einer systematischen Darstellung der Analyse nichtlinearer dynamische Netzwerke*, FB: Elektrotechnik, Universität Kaiserslautern, Dissertation.
- [16] **Houwen, P.J. van der & Sommeijer, B.P.** (1990): *Parallel iteration of high-order Runge-Kutta methods with stepsize control*, J. Comput. Appl. Math. **29**, pp. 111-127.
- [17] **Houwen, P.J. van der, Sommeijer, B.P. & Couzy, W.** (1992): *Embedded diagonally implicit Runge-Kutta algorithms on parallel computers*, to appear in Math. Comp.
- [18] **Houwen, P.J. van der & Sommeijer, B.P.** (1991): *Iterated Runge-Kutta methods on parallel computers*, SIAM J. Sci. Stat. Comput. **12**, pp. 1000-1028.
- [19] **Houwen, P.J. van der & Sommeijer, B.P.** (1991): *Analysis of parallel diagonal-implicit iteration of Runge-Kutta methods*, Report R-NM9114, Centre for Mathematics and Computer Science, Amsterdam, submitted for publication.
- [20] **Iserles, A. & Nørsett, S.P.** (1991): *On the theory of parallel Runge-Kutta methods*, IMA J. Numer. Anal. **10**, pp. 463-488.
- [21] **Jackson, K. & Nørsett, S.P.** (1990): *The potential for parallelism in Runge-Kutta methods, Part I: RK formulas in standard form*, Technical Report No. 239/90, Dept. of Computer Science, University of Toronto.
- [22] **Jackson, K. & Nørsett, S.P.**: *The potential for parallelism in Runge-Kutta methods, Part II: RK predictor-corrector formulas*, in preparation.
- [23] **Lie, I.** (1987): *Some aspects of parallel Runge-Kutta methods*, Report No. 3/87, Dept. of Mathematics, University of Trondheim.
- [24] **Nørsett, S.P.** (1974): *Semi-explicit Runge-Kutta methods*, Report Mathematics and Computation No. 6/74, Dept. of Mathematics, University of Trondheim.
- [25] **Nørsett, S.P. & Thomsen, P.G.** (1984): *Embedded SDIRK-methods of basic order three*, BIT **24**, pp. 634-646.
- [26] **Nørsett, S.P. & Simonsen, H.H.** (1989): *Aspects of parallel Runge-Kutta methods*, in: *Numerical methods for ordinary differential equations*, A. Bellen, C.W. Gear & E. Russo (eds.), Proceedings of the L'Aquila conference, 1987, Lecture Notes in Mathematics, Vol. **1386**, Springer-Verlag, Berlin, pp. 103-117.
- [27] **Prothero, A. & Robinson, A.** (1974): *On the stability and accuracy of one-step methods for solving stiff systems of ordinary differential equations*, Math. Comp. **28**, pp. 145-162.
- [28] **Shampine, L.F.** (1980): *Implementation of implicit formulas for the solution of ODEs*, SIAM J. Sci. Stat. Comput. **1**, pp. 103-118.
- [29] **Wolfbrandt, A.** (1977): *A study of Rosenbrock processes with respect to order conditions and stiff stability*, Thesis, Chalmers University of Technology, Göteborg, Sweden.

