

Algorithms and Data Structures for Multi-Adaptive Time-Stepping

JOHAN JANSSON

Royal Institute of Technology, Stockholm

and

ANDERS LOGG

University of Oslo

Multi-adaptive Galerkin methods are extensions of the standard continuous and discontinuous Galerkin methods for the numerical solution of initial value problems for ordinary or partial differential equations. In particular, the multi-adaptive methods allow individual and adaptive time steps to be used for different components or in different regions of space. We present algorithms for efficient multi-adaptive time-stepping, including the recursive construction of time slabs and adaptive time step selection. We also present data structures for efficient storage and interpolation of the multi-adaptive solution. The efficiency of the proposed algorithms and data structures is demonstrated for a series of benchmark problems.

Categories and Subject Descriptors: G.1.7 [**Numerical Analysis**]: Ordinary Differential Equations—*Error analysis; Initial value problems*; G.1.8 [**Numerical Analysis**]: Partial Differential Equations—*Finite element methods*; G.4 [**Mathematical Software**]:—*Algorithm design and analysis; Efficiency*

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Multi-adaptivity, individual time steps, local time steps, multirate, ODE, continuous Galerkin, discontinuous Galerkin, mcgq, mdgq, C++, implementation, algorithms, DOLFIN

ACM Reference Format:

Jansson, J. and Logg, A. 2008. Algorithms and data structures for multi-adaptive time-stepping. *ACM Trans. Math. Softw.* 35, 3, Article 17 (October 2008), 24 pages. DOI = 10.1145/1391989.1391990. <http://doi.acm.org/10.1145/1391989.1391990>.

A. Logg is supported by an Outstanding Young Investigator grant from the Research Council of Norway, NFR 180450.

Authors' addresses: J. Jansson, School of Computer Science and Communication, Royal Institute of Technology, SE-100 44 Stockholm, Sweden; email: jjan@csc.kth.se; A. Logg, Center for Biomedical Computing, Simula Research Laboratory, P.O. Box 134, 1325 Lysaker, Norway; email: logg@simula.no.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credits is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from the Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2008 ACM 0098-3500/2008/10-ART17 \$5.00 DOI: 10.1145/1391989.1391990. <http://doi.acm.org/10.1145/1391989.1391990>.

ACM Transactions on Mathematical Software, Vol. 35, No. 3, Article 17, Pub. date: October 2008.

1. INTRODUCTION

In an earlier sequence of papers [Logg 2003a; 2003b; 2006], we introduced the multi-adaptive Galerkin methods mcG(q) and mdG(q) for the approximate numerical solution of ODEs of the form:

$$\begin{aligned} \dot{u}(t) &= f(u(t), t), \quad t \in (0, T], \\ u(0) &= u_0, \end{aligned} \tag{1}$$

where $u : [0, T] \rightarrow \mathbb{R}^N$ is the solution to be computed, $u_0 \in \mathbb{R}^N$ is a given initial value, $T > 0$ is a given final time, and $f : \mathbb{R}^N \times (0, T] \rightarrow \mathbb{R}^N$ is a given function that is Lipschitz continuous in u and bounded.

The multi-adaptive Galerkin methods mcG(q) and mdG(q) extend the standard mono-adaptive continuous and discontinuous Galerkin methods cG(q) and dG(q), studied previously in Hulme [1972b, 1972a]; Jamet [1978]; Delfour et al. [1981]; Eriksson et al. [1985]; Johnson [1988]; Eriksson and Johnson [1991, 1995a, 1995b, 1995c]; Eriksson et al. [1998, 1995]; Estep [1995]; Estep and French [1994]; Estep et al. [2000]; Estep and Williams [1996] and Estep and Stuart [2002], by allowing individual time step sequences $k_i = k_i(t)$ for the different components $U_i = U_i(t)$, $i = 1, 2, \dots, N$, of the approximate solution $U \approx u$ of the initial value problem (1). For related work on local time-stepping, see also Hughes et al. [1983a, 1983b]; Makino and Aarseth [1992]; Davé et al. [1997]; Alexander and Agnor [1998]; Osher and Sanders [1983]; Flaherty et al. [1997]; Dawson and Kirby [2001]; Lew et al. [2003]; Engstler and Lubich [1997]; Savcenco et al. [2005] and Savcenco [2008]. In comparison with existing methods for local time-stepping, the main advantage of the multi-adaptive Galerkin methods mcG(q) and mdG(q) is the automatic local step size selection based on a global a posteriori error estimate built into these methods.

In the current article, we discuss important aspects of the implementation of multi-adaptive Galerkin methods. While earlier results on multi-adaptive time-stepping presented in Logg [2003a, 2003b, 2006] include the formulation of the methods, a priori and a posteriori error estimates, together with a proof-of-concept implementation and results for a number of model problems, the current article addresses the important issue of efficiently implementing the multi-adaptive methods with minimal overhead as compared to standard mono-adaptive solvers. For many problems, in particular when the propagation of the solution is local in space and time, the potential speedup of multi-adaptivity is large, but the actual speedup may be far from the ideal speedup if the overhead of the more complex implementation is significant.

1.1 Implementation

The algorithms presented in this article are implemented by the multi-adaptive ODE-solver available in DOLFIN [Logg et al. <http://www.fenics.org/dolfin/>; Hoffman and Logg 2002], Dynamic Object-oriented Library for FINite element computation, which is the C++ interface of the new open-source software project FEniCS [FEniCS 2008; Logg 2007; Dupont et al. 2003] for the automation of Computational Mathematical Modeling (CMM). The multi-adaptive solver in DOLFIN is based on the original implementation

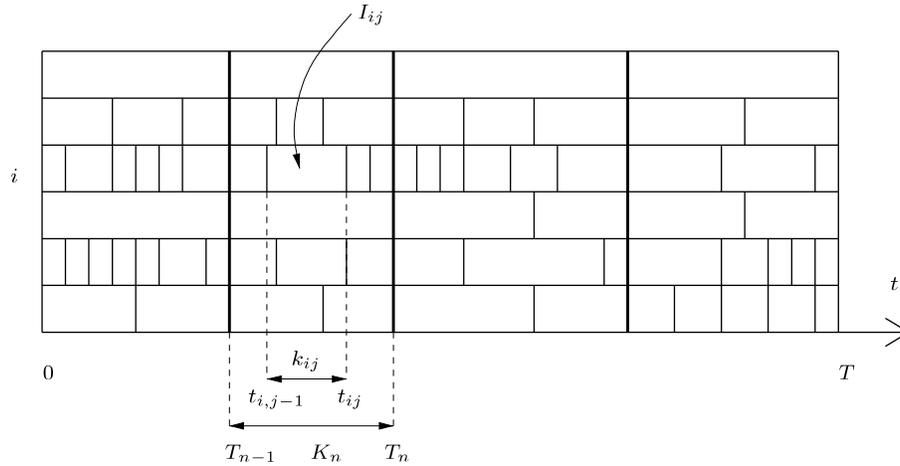


Fig. 1. Individual partitions of the interval $(0, T]$ for different components. Elements between common synchronized time levels are organized in time slabs. In this example, we have $N = 6$ and $M = 4$.

Tanganyika, presented in Logg [2003b], but has been completely rewritten for DOLFIN and is being actively developed by the authors.

1.2 Obtaining the Software

DOLFIN is licensed under the GNU (Lesser) General Public License [Free Software Foundation 1999], which means that anyone is free to use or modify the software, provided these rights are preserved. The complete source code of DOLFIN, including numerous example programs, is available at the DOLFIN Web page [Logg et al. <http://www.fenics.org/dolfin/>].

1.3 Notation

The following notation is used throughout this article: each component $U_i(t)$, $i = 1, \dots, N$, of the approximate $m(c/d)G(q)$ solution $U(t)$ of (1) is a piecewise polynomial on a partition of $(0, T]$ into m_i sub-intervals. Sub-interval j for component i is denoted by $I_{ij} = (t_{i,j-1}, t_{ij}]$, and the length of the sub-interval is given by the local *time step* $k_{ij} = t_{ij} - t_{i,j-1}$. We shall sometimes refer to I_{ij} as an *element*. This is illustrated in Figure 1. On each sub-interval I_{ij} , $U_i|_{I_{ij}}$ is a polynomial of degree at most q_{ij} .

Furthermore, we shall assume that the interval $(0, T]$ is partitioned into blocks between certain synchronized time levels $0 = T_0 < T_1 < \dots < T_M = T$. For each T_n , $n = 0, 1, \dots, M$ and each $i = 1, 2, \dots, N$, we require that there is a $0 \leq j \leq m_i$ such that $t_{ij} = T_n$. We refer to the collection of local intervals between two synchronized time levels T_{n-1} and T_n as a *time slab*. We denote the length of a time slab by $K_n = T_n - T_{n-1}$.

1.4 Outline of the Article

We first give an introduction to multi-adaptive time-stepping in Section 2. We then present the key algorithms used by the multi-adaptive ODE solver

of DOLFIN in Section 3, followed by a discussion of data structures for efficient representation and interpolation of multi-adaptive solutions in Section 4. In Section 5, we discuss the efficiency of multi-adaptive time-stepping and in Section 6, we present a number of numerical examples that demonstrate the efficiency of the proposed algorithms and data structures. Finally, we give some concluding remarks in Section 7.

2. MULTI-ADAPTIVE TIME-STEPPING

In this section, we give a quick introduction to multi-adaptive time-stepping, including the formulation of the methods, error estimates, and adaptivity. For a more detailed account, we refer the reader to Logg [2003a, 2003b, 2006].

2.1 Formulation of the Methods

The mcG(q) and mdG(q) methods are obtained by multiplying the system of equations (1) with a suitable test function v , to obtain the following variational problem. Find $U \in V$ with $U(0) = u_0$, such that:

$$\int_0^T (v, \dot{U}) dt = \int_0^T (v, f(U, \cdot)) dt \quad \forall v \in \hat{V}, \quad (2)$$

where (\cdot, \cdot) denotes the standard inner product on \mathbb{R}^N and (\hat{V}, V) is a suitable pair of discrete function spaces, the *test* and *trial* spaces respectively.

For the standard cG(q) method, the trial space V consists of the space of continuous piecewise polynomial vector-valued functions of degree $q = q(t)$ on a partition $0 = t_0 < t_1 < \dots < t_M = T$ and the test space \hat{V} consists of the space of (possibly discontinuous) piecewise polynomial vector-valued functions of degree $q - 1$ on the same partition. The multi-adaptive mcG(q) method extends the standard cG(q) method by extending the test and trial spaces to piecewise polynomial spaces on individual partitions of the time interval that satisfy the constraints introduced in the previous section and illustrated in Figure 1. Thus, each component $U_i = U_i(t)$ is continuous and a piecewise polynomial on the individual partition $0 = t_{i0} < t_{i1} < \dots < t_{im_i} = T$ for $i = 1, 2, \dots, N$.

For the standard dG(q) method, the test and trial spaces are equal and consist of the space of (possibly discontinuous) piecewise polynomial vector-valued functions of degree $q = q(t)$ on a partition $0 = t_0 < t_1 < \dots < t_M = T$, which extends naturally to the multi-adaptive mdG(q) method by allowing each component of the test and trial functions to be a piecewise polynomial on its own partition of the time interval as shown for the mcG(q) method. Note that for both the dG(q) and the mdG(q) methods, the integral $\int_{0,T} (v, \dot{U}) dt$ in (2) must be treated appropriately at the points of discontinuity, see Logg [2003a].

Both in the case of the mcG(q) and mdG(q) methods, the variational problem (2) gives rise to a system of discrete equations by expanding the solution U in a suitable basis on each local interval I_{ij} ,

$$U_i|_{I_{ij}} = \sum_{m=0}^{q_{ij}} \xi_{ijm} \phi_{ijm}, \quad (3)$$

where $\{\xi_{ijm}\}_{m=0}^{q_{ij}}$ are the *degrees of freedom* for U_i on I_{ij} and $\{\phi_{ijm}\}_{m=0}^{q_{ij}}$ is a suitable basis for $P^{q_{ij}}(I_{ij})$. For any particular choice of quadrature, the resulting system of discrete equations takes the form of an implicit Runge–Kutta method on each local interval I_{ij} . The discrete equations take the form:

$$\xi_{ijm} = \xi_{ij0}^- + k_{ij} \sum_{n=0}^{q_{ij}} w_{mn}^{[q_{ij}]} f_i(U(\tau_{ij}^{-1}(s_n^{[q_{ij}]}), \tau_{ij}^{-1}(s_n^{[q_{ij}]}), \quad (4)$$

for $m = 0, \dots, q_{ij}$, where $\{w_{mn}^{[q_{ij}]}\}_{m=0, n=0}^{q_{ij}}$ are weights, τ_{ij} maps I_{ij} to $(0, 1]$, $\tau_{ij}(t) = (t - t_{i,j-1}) / (t_{ij} - t_{i,j-1})$, and $\{s_n^{[q_{ij}]}\}_{n=0}^{q_{ij}}$ are quadrature points defined on $[0, 1]$. Note that we have assumed that the number of quadrature points is equal to the number of nodal points. See Logg [2003a] for a discussion of suitable quadrature rules and basis functions.

2.2 Error Estimates and Adaptivity

The global error $e = U - u$ of the approximate solution U of (1) may be bounded in terms of computable quantities. Such an a posteriori error estimate is proved in Logg [2003a], both for the mcG(q) and mdG(q) methods. The a posteriori error estimate provides a bound for any given linear functional $\mathcal{M} : \mathbb{R}^N \rightarrow \mathbb{R}$ of the global error $e(T)$ at the final time, such as the error $e_i(T)$, in a single component. Bounds for the error itself in various norms may also be approximated. In the following, we state the basic a posteriori error estimate for the mcG(q) method and refer to Logg [2003a] for a complete discussion, including error estimates for mdG(q).

For the mcG(q) method, the error estimate takes the following form:

$$|\mathcal{M}(e(T))| \leq E \equiv \sum_{i=1}^N S_i(T) \max_{[0, T]} \{C_i k_i^{q_i} |R_i|\}. \quad (5)$$

Here, $R = \dot{U} - f(U, \cdot)$ denotes the *residual* of the computed solution, $C_i = C_i(t)$ denotes an interpolation constant (which may be different for each local interval) and $S_i(T)$ denotes a *stability factor* that measures the rate of propagation of local errors for component U_i (the influence of a nonzero residual in component U_i on the size of the error in the given functional). By selecting the local time steps $k_i = k_i(t)$ such that $E = \text{TOL}$ for a given tolerance TOL, one may thus guarantee that the error in the functional \mathcal{M} is bounded by the given tolerance, $|\mathcal{M}(e(T))| \leq \text{TOL}$.

This can be compared to standard Runge–Kutta methods for the solution of initial value problems, where the stability factor quantifies the relationship between the local error and the global error. Note that alternatively, the stability information may be kept as a local time-dependent *stability weight* for more fine-grained control of the contributions to the global error. The stability factors are obtained by solving a dual problem of (1) for the given functional \mathcal{M} , [Eriksson et al. 1995; Logg 2003a]. The particular form of the dual problem for (1) will be discussed in Section 3.5.

The individual time steps may be chosen so as to equidistribute the error in the different components in an attempt to satisfy

$$C_{ij} k_{ij}^{q_{ij}} \max_{I_{ij}} |R_i| = \text{TOL}/(NS_i(T)), \quad (6)$$

for each local time interval I_{ij} . This may be done in an iterative fashion, as outlined in the following basic adaptive algorithm:

- (0) Assume $S_i(T) = 1$ for $i = 1, 2, \dots, N$;
- (i) solve the primal problem with time steps based on (6);
- (ii) solve the dual problem and compute the stability factors;
- (iii) compute an error bound E based on (5);
- (iv) if $E \leq \text{TOL}$ then stop; if not go back to (i).

3. ALGORITHMS

We present a collection of key algorithms for multi-adaptive time-stepping. The algorithms are given in pseudo-code and where appropriate we give remarks on how the algorithms have been implemented in C++ for DOLFIN. In most cases, we present simplified versions of the algorithms with a focus on the most essential steps.

3.1 General Algorithm

The general multi-adaptive time-stepping algorithm is Algorithm 1. Starting at $t = 0$, the algorithm creates a sequence of time slabs until the given end time T , is reached. In each macro time step, Algorithm 2 (CreateTimeSlab) is called to create a time slab covering an interval $[T_{n-1}, T_n]$ such that $T_n \leq T$. For each time slab, the system of discrete equations is solved iteratively, using direct fixed-point iteration or a preconditioned Newton's method, until the discrete equations given by the mcG(q) or mdG(q) method have converged.

Algorithm 1. $U = \text{Integrate}(\text{ODE})$

```

t ← 0
while t < T
  {time slab, t} ← CreateTimeSlab({1, ..., N}, t, T)
  SolveTimeSlab(time slab)
end while

```

The basic forward integrator, Algorithm 1, can be used as the main component of an adaptive algorithm with automated error control of the computed solution, as outlined in Section 2. In each iteration, the *primal problem* (1) is solved using Algorithm 1. An ODE of the form (1) representing the *dual problem* is then created and solved using Algorithm 1. It is important to note that both the primal and the dual problems may be solved using the same algorithm, but with possibly different time steps, tolerances, methods, and orders. When the solution of the dual problem has been computed, the stability factors $\{S_i(T)\}_{i=1}^N$, and the error estimate may be computed.

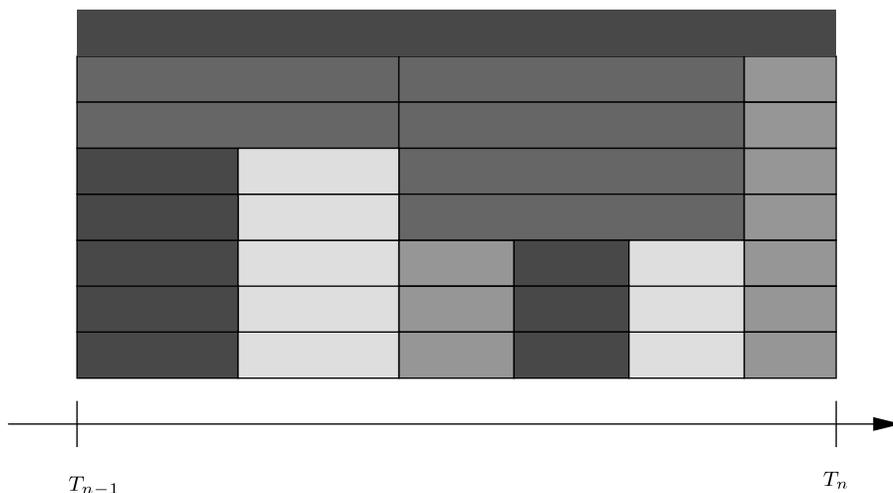


Fig. 2. The recursive organization of the time slab. Each time slab contains an element group and a list of recursively nested time slabs. The root time slab in the figure contains one element group of one element and three sub-slabs. The first of these sub-slabs contains an element group of two elements and two nested sub-slabs, and so on. The root time slab recursively contains a total of nine element groups and 33 elements.

3.2 Recursive Construction of Time Slabs

In each step of Algorithm 1, a new time slab is created between two synchronized time levels T_{n-1} and T_n . The time slab is organized recursively as follows. The root time slab covering the interval $[T_{n-1}, T_n]$ contains a non-empty list of elements, which we refer to as an *element group*, and a possibly empty list of time slabs, which in turn may contain nested groups of elements and time slabs. Each such element group together with the corresponding nested set of element groups is referred to as a *sub-slab*. This is illustrated in Figure 2.

To create a time slab, we first compute the desired time steps for all components as given by the a posteriori error estimate (5). We discuss the time step selection in detail in Section 3.4. A threshold θK , is then computed based on the maximum time step K , among the components and a fixed parameter $\theta \in (0, 1)$, controlling the density of the time slab. The components are partitioned into two sets based on the threshold, and a large time step \underline{K} is selected to be the smallest time step among the components in the set with large time steps as described in Algorithm 3 and illustrated in Figure 3. For each component in the group with large time steps, an element is created and added to the element group of the time slab. The remaining components, with small time steps, are processed by a recursive application of this algorithm for the construction of time slabs.

We organize the recursive construction of time slabs as described by Algorithms 2, 3, 4, and 5. The recursive construction simplifies the implementation; each recursively nested sub-slab can be considered as a sub-system of the ODE. Note that the element group containing elements for components in group I_1 is created before the recursively nested sub-slabs for components in

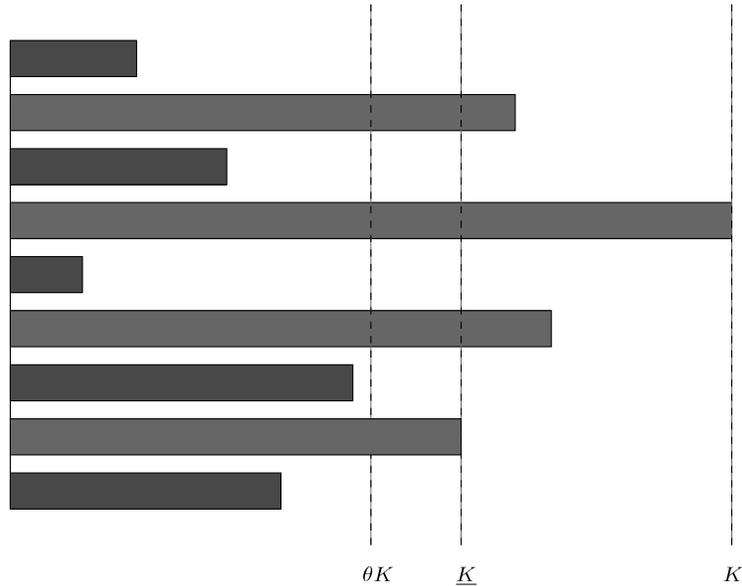


Fig. 3. The partition of components into groups of small and large time steps for $\theta = 1/2$.

group I_0 . The tree of time slabs is thus created recursively *breadth-first*, which means in particular that the element for the component with the largest time step is created first.

Algorithm 3 for the partition of components can be implemented efficiently using the function `std::partition()`, which is part of the Standard C++ Library.

3.3 Solving the System of Discrete Equations

On each time slab \mathcal{T}_n , $n = 1, 2, \dots, M$, we need to solve a system of equations for the degrees of freedom on the time slab. On each local interval, $I_{ij} \in \mathcal{T}_n$, these equations are given by (4). Depending on the properties of the given system (1), different solution strategies for the time slab system (4), may be appropriate as outlined in the following.

Algorithm 2. $\{\text{time slab}, T_n\} = \text{CreateTimeSlab}(\text{components}, T_{n-1}, T)$

$\{I_0, I_1, K\} \leftarrow \text{Partition}(\text{components})$

if $T_{n-1} + K < T$

$T_n \leftarrow T_{n-1} + K$

else

$T_n \leftarrow T$

end if

element group $\leftarrow \text{CreateElements}(I_1, T_{n-1}, T_n)$

time slabs $\leftarrow \text{CreateTimeSlabs}(I_0, T_{n-1}, T_n)$

time slab $\leftarrow \{\text{element group}, \text{time slabs}\}$

Algorithm 3. $\{I_0, I_1, K\} = \text{Partition}(\text{components})$

```

 $I_0 \leftarrow \emptyset$ 
 $I_1 \leftarrow \emptyset$ 
 $K \leftarrow \text{maximum time step within components}$ 
for each component
   $k \leftarrow \text{time step of component}$ 
  if  $k < \theta K$ 
     $I_0 \leftarrow I_0 \cup \{\text{component}\}$ 
  else
     $I_1 \leftarrow I_1 \cup \{\text{component}\}$ 
  endif
end for
 $\underline{K} \leftarrow \text{minimum time step within } I_1$ 
 $K \leftarrow \underline{K}$ 

```

Algorithm 4. $\text{elements} = \text{CreateElements}(\text{components}, T_{n-1}, T_n)$

```

 $\text{elements} \leftarrow \emptyset$ 
for each component
   $\text{create element for component on } [T_{n-1}, T_n]$ 
   $\text{elements} \leftarrow \text{elements} \cup \text{element}$ 
end for

```

Algorithm 5. $\text{time slabs} = \text{CreateTimeSlabs}(\text{components}, T_{n-1}, T_n)$

```

 $\text{time slabs} \leftarrow \emptyset$ 
 $t \leftarrow T_{n-1}$ 
while  $t < T$ 
   $\{\text{time slab}, t\} \leftarrow \text{CreateTimeSlab}(\text{components}, t, T_n)$ 
   $\text{time slabs} \leftarrow \text{time slabs} \cup \text{time slab}$ 
end while

```

3.3.1 *Direct Fixed-Point Iteration.* In the simplest case, the time slab system is solved by direct fixed-point iteration on (4) for each element in the time slab. The fixed-point iteration is performed in a forward fashion, sweeping over the elements in the time slab in the same order as they are created by Algorithm 2. This means that for each component in the time slab system, the end-time value on each element is updated before the degrees of freedom for the following element. Thus, for each element $I_{ij} \in \mathcal{T}_n$, we compute the degrees of freedom $\{\xi_{ijm}\}_{j=0}^{q_{ij}}$ according to:

$$\xi_{ijm} = \xi_{ij0}^- + k_{ij} \sum_{n=0}^{q_{ij}} w_{mn}^{[q_{ij}]} f_i(U(\tau_{ij}^{-1}(s_n^{[q_{ij}]}), \tau_{ij}^{-1}(s_n^{[q_{ij}]}), \quad m = 0, 1, \dots, q_{ij}. \quad (7)$$

Direct fixed-point iteration converges if the system is non-stiff and typically only a few iterations are needed. In fact, one may consider a system to be stiff if direct fixed-point iteration does not converge.

3.3.2 Damped Fixed-Point Iteration. If the system is stiff, that is, direct fixed-point iteration does not converge, one may introduce a suitable amount of damping to adaptively stabilize the fixed-point iteration. The fixed-point iteration (7) may be written in the form:

$$\xi_{ijm} = g_{ijm}(\xi), \quad (8)$$

where ξ is the vector of degrees of freedom for the solution on the time slab. We modify the fixed-point iteration by introducing a damping parameter α :

$$\xi_{ijm} = (1 - \alpha_{ijm})\xi_{ijm} + \alpha_{ijm}g_{ijm}(\xi). \quad (9)$$

In Logg [2004], a number of different strategies for the selection of the damping parameter α are discussed. We mention two of these strategies here. The first strategy chooses α based on the diagonal derivatives $\partial f_i / \partial u_i$, $i = 1, 2, \dots, N$, corresponding to a modified Newton's method where the Jacobian is approximated by a diagonal matrix. This strategy works well for systems with a diagonally dominant Jacobian, including many systems arising when modeling chemical reactions. The second strategy adaptively chooses a scalar α , based on the convergence of the fixed-point iterations.

3.3.3 Newton's Method. Alternatively, one may apply Newton's method directly to the full system of equations (7) associated with each time slab. The linear system in each Newton iteration may then be solved either by a direct method or an iterative method such as a Krylov subspace method in combination with a suitable preconditioner, depending on the characteristics of the underlying system (1). In addition, one may also apply a special preconditioner that improves the convergence by propagating values forward in time within the time slab. Note that if the multi-adaptive efficiency index is large (see Section 5), then the time slab system is not significantly larger than the corresponding time slab system for a mono-adaptive method.

3.3.4 Choosing a Solution Strategy. Ultimately, an intelligent solver should automatically choose a suitable algorithm for the solution of the time slab system. Thus, the solver may initially try direct fixed-point iteration. If the system is stiff, the solver switches to adaptive fixed-point iteration (as outlined in Logg [2004]). Finally, if the adaptive fixed-point iteration converges slowly, the solver may switch to Newton's method.

3.3.5 Interpolation of the Solution. To update the degrees of freedom on an element according to (7), the appropriate component f_i of the right-hand side of (1) needs to be evaluated at the set of quadrature points. In order for f_i to be evaluated, each component $U_{i'}$ of the computed solution U on which f_i depends, needs to be evaluated at the quadrature points. We let $S_i \subseteq \{1, \dots, N\}$ denote the *sparsity pattern* of component U_i , that is, the set of components on which f_i depends,

$$S_i = \{i' \in \{1, \dots, N\} : \partial f_i / \partial u_{i'} \neq 0\}. \quad (10)$$

Thus, to evaluate f_i at a given quadrature point t , only the components $\{U_{i'}\}_{i' \in S_i}$ need to be evaluated at t , as in Algorithm 6. This is of particular importance for

problems of sparse structure and enables efficient multi-adaptive integration of time-dependent PDEs, as demonstrated in Section 6. The sparsity pattern \mathcal{S}_i , is automatically detected by the solver. Alternatively, the sparsity pattern may be specified by a sparse matrix.

Algorithm 6. $y = \text{EvaluateRightHandSide}(i, t)$

```

for  $i' \in \mathcal{S}_i$ 
     $x(i') \leftarrow U_{i'}(t)$ 
end for
 $y \leftarrow f_i(x, t)$ 

```

In Algorithm 6, the key step is the evaluation of a component $U_{i'}$ at a given point t . For a standard mono-adaptive method, this is straightforward, since all components use the same time steps. In particular, if the quadrature points are chosen to be the same as the nodal points, the value of $U_{i'}(t)$ is known. For a multi-adaptive method, a quadrature point t for the evaluation of f_i is not necessarily a nodal point for $U_{i'}$. To evaluate $U_{i'}(t)$, one thus needs to find the local interval $I_{i'j}$, such that $t \in I_{i'j}$ and then evaluate $U_{i'}(t)$ by interpolation on that interval. In Section 4, we discuss data structures that allow efficient storage and interpolation of the multi-adaptive solution. These data structures give $\mathcal{O}(1)$ access to the value of any component $U_{i'}$ in the sparsity pattern \mathcal{S}_i at any quadrature point t for f_i .

3.4 Multi-Adaptive Time Step Selection

The individual and adaptive time steps k_{ij} , are determined during the recursive construction of time slabs based on an a posteriori error estimate, as discussed in Section 2. Thus, according to (6), each local time step k_{ij} should be chosen to satisfy:

$$k_{ij} = \left(\frac{\text{TOL}}{C_{ij} N \mathcal{S}_i(T) \max_{I_{ij}} |R_i|} \right)^{1/q_{ij}}, \quad (11)$$

where TOL is a given tolerance.

However, the time steps cannot be based directly on (11), since that leads to unwanted oscillations in the size of the time steps. If $r_{i,j-1} = \max_{I_{i,j-1}} |R_i|$ is small, then k_{ij} will be large, and as a result r_{ij} will also be large. Consequently, $k_{i,j+1}$ and $r_{i,j+1}$ will be small, and so on. To avoid these oscillations, we adjust the time step k_{ij} according to Algorithm 7, which determines the new time step as a weighted harmonic mean value of the previous time step and the time step given by (11). Alternatively, DOLFIN provides time step control based on the PID controllers presented in Gustafsson et al. [1988] and Söderlind [2003], including H0211 and H211PI. However, the simple controller of Algorithm 7 performs well compared to the more sophisticated controllers in Gustafsson et al. [1988] and Söderlind [2003]. A suitable value for the weight w , in Algorithm 7 is $w = 5$ (found empirically).

The initial time steps $k_{11} = k_{21} = \dots = k_{N1} = K_1$ are chosen equal for all components and are determined iteratively for the first time slab. The size K_1

Algorithm 7. $k = \text{Controller}(k_{\text{new}}, k_{\text{old}}, k_{\text{max}})$

$k \leftarrow (1 + w)k_{\text{old}}k_{\text{new}}/(k_{\text{old}} + wk_{\text{new}})$

$k \leftarrow \min(k, k_{\text{max}})$

of the first time slab is first initialized to some default value, possibly based on the length T of the time interval, and then adjusted until the local residuals are sufficiently small for all components.

3.5 Solving the Dual Problem

Stability factors may be approximated by numerically solving an auxiliary dual problem for (1). This dual problem is given by the following system of linear ordinary differential equations:

$$\begin{aligned} -\dot{\varphi}(t) &= J(U(t), t)^\top \varphi(t), \quad t \in [0, T), \\ \varphi(T) &= \psi, \end{aligned} \tag{12}$$

where $J(U(t), t)$ denotes the Jacobian of the right-hand side f , of (1), at time t and $\psi = \mathcal{M}'$ (the Riesz representer of \mathcal{M}) is initial data for the dual problem corresponding to the given functional \mathcal{M} to be estimated. Note that we need to linearize around the computed solution U , since the exact solution u , of (1), is not known. To solve this backward problem over $[0, T)$ using the forward integrator Algorithm 1, we rewrite (12) as a forward problem. With $w(t) = \varphi(T - t)$, we have $\dot{w} = -\dot{\varphi}(T - t) = J(U(T - t), T - t)^\top w(t)$, and so (12) can be written as a forward problem for w in the form:

$$\begin{aligned} \dot{w}(t) &= f^*(w(t), t) \equiv J(U(T - t), T - t)^\top w(t), \quad t \in (0, T], \\ w(0) &= \psi. \end{aligned} \tag{13}$$

4. DATA STRUCTURES

For a standard mono-adaptive method, the solution on a time slab is typically stored as an array of values at the right end-point of the time slab, or as a list of arrays (possibly stored as one contiguous array) for a higher order method with several stages. However, a different data structure is needed to store the solution on a multi-adaptive time slab. Such a data structure should ideally store the solution with minimal overhead compared to the cost of storing only the array of degrees of freedom for the solution on the time slab. In addition, it should also allow for efficient interpolation of the solution, that is, accessing the values of the solution for all components at any given time within the time slab. We present a data structure that allows efficient storage of the entire solution on a time slab with little overhead, and at the same time allows efficient interpolation with $\mathcal{O}(1)$ access to any given value during the iterative solution of the system of discrete equations.

4.1 Representing the Solution

The multi-adaptive solution on a time-slab can be efficiently represented using a data structure consisting of eight arrays, as shown in Table I. For simplicity,

Table I. Data Structures for Efficient Representation of a Multi-Adaptive Time Slab

Array	Type	Description
sa	double	left end-points for sub-slabs
sb	double	right end-points for sub-slabs
jx	double	values for degrees of freedom
ei	int	component indices for elements
es	int	time slabs containing elements
ee	int	previous elements for elements
ed	int	first dependencies for elements
de	int	elements for dependencies

we assume that all elements in a time slab are constructed for the same choice of method, $\text{mcG}(q)$ or $\text{mdG}(q)$, for a given fixed q .

The recursive construction of time slabs as discussed in Section 3.2 generates a sequence of *sub slabs*, each containing a list of *elements* (an element group). For each sub-slab, we store the value of the time t at the left end-point and at the right end-point in the two arrays sa and sb . Thus, for sub-slab number s covering the interval (a_s, b_s) , we have:

$$\begin{aligned} a_s &= \text{sa}[s], \\ b_s &= \text{sb}[s]. \end{aligned} \quad (14)$$

Furthermore, for all elements in the root time slab, we store the degrees of freedom in the order they are created in the array jx (mapping a degree of freedom j , to the value x , of that degree of freedom). Thus, if each element has q degrees of freedom, as in the case of the multi-adaptive $\text{mcG}(q)$ method, then the length of the array jx , is q times the number of elements. In particular, if all components use the same time steps, then the length of the array jx , is qN .

For each element, we store the corresponding component index i , in the array ei , in order to be able to evaluate the correct component f_i , of the right-hand side f , of (1) when iterating over all elements in the time slab to update the degrees of freedom. When updating the values on an element according to (7), it is also necessary to know the left and right end-points of the elements. Thus, we store an array es , that maps the number e , of a given element to the number s , of the corresponding sub-slab containing the element. As a consequence, the left end-point a_e , and right end-point b_e , for a given element e , are given by:

$$\begin{aligned} a_e &= \text{sa}[\text{es}[e]], \\ b_e &= \text{sb}[\text{es}[e]]. \end{aligned} \quad (15)$$

4.2 Interpolating the Solution at Quadrature Points

Updating the values on an element according to (7) also requires knowledge of the value at the left end-point, which is given as the end-time value on the previous element in the time slab for the same component (or the end-time value from the previous time slab). This information is available in the array ee , which stores for each element, the number of the previous element (or -1 if there is no previous element).

As discussed in Section 3.3, the system of discrete equations on each time slab is solved by iterating over the elements in the time slab and updating the values on each element, either in a direct fixed-point iteration or a Newton's method. We must then, for any given element e , corresponding to some component $i = \text{ei}[e]$, evaluate the right-hand side f_i , at each quadrature point t , within the element. This requires the values of the solution U , at t , for all components contained in the sparsity pattern S_i , for component i , according to Algorithm 6. As a consequence of Algorithm 2 for the recursive construction of time slabs, elements for components that use large time steps are constructed before elements for components that use small time steps. Since all elements of the time slab are traversed in the same order during the iterative solution of the system of discrete equations, elements corresponding to large time steps have recently been visited and cover any element that corresponds to a smaller time step. The last visited element for each component is stored in an auxiliary array `elast`, of size N . Thus, if $i' \in S_i$, and component i' , has recently been visited, then it is straightforward to find the latest element $e' = \text{elast}[i']$, for component i' , that covers the current element for component i , and interpolate $U_{i'}$, at time t . It is also straightforward to interpolate the values for any components that are present in the same element group as the current element.

However, when updating the values on an element e , corresponding to some component $i = \text{ei}[e]$, depending on some other component $i' \in S_i$, which uses smaller time steps, one must find for each quadrature point t , on the element e , the element e' , for component i' , containing t , which is nontrivial. The element e' , can be found by searching through all elements for component i' , in the time slab, but this quickly becomes inefficient. Instead, we store for each element e , a list of dependencies to elements with smaller time steps in the two arrays `ed`, and `de`. These two arrays store a sparse integer matrix of dependencies to elements with smaller time steps for all elements in the time slab. Thus, for any given element e , the number of dependencies to elements with smaller time steps is given by:

$$\text{ed}[e + 1] - \text{ed}[e], \quad (16)$$

and the elements with smaller time steps that need to be interpolated at the quadrature points for element e , are given by:

$$\{\text{de}[\text{ed}[e]], \text{de}[\text{ed}[e] + 1], \dots, \text{de}[\text{ed}[e + 1] - 1]\}. \quad (17)$$

5. PERFORMANCE

The efficiency of multi-adaptive time-stepping compared to standard mono-adaptive time-stepping depends on the system being integrated, the tolerance, and the efficiency of the implementation. For many systems, the potential speedup is large, but the actual speedup also depends on the overhead needed to handle the additional complications of a multi-adaptive implementation: the recursive construction of time slabs and the interpolation of values within a time slab.

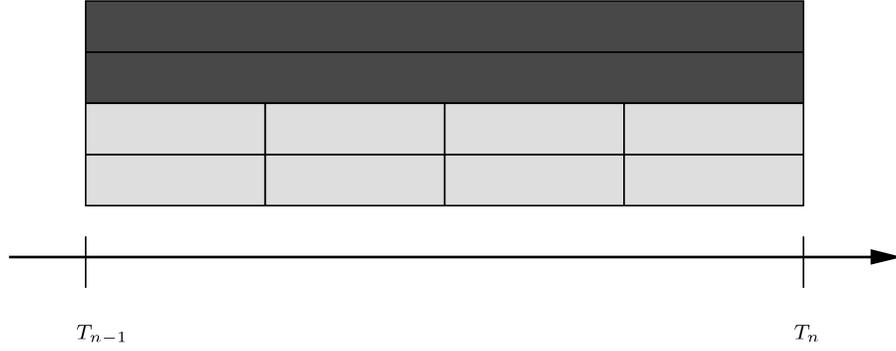


Fig. 4. A time slab with $N_K = N_k = 2$ and multi-adaptive efficiency index $\mu = 16/10 = 1.6$.

To study the performance of multi-adaptive time-stepping, we consider a system of N components and time steps given by $\{k_{ij} = |I_{ij}| : I_{ij} \in \mathcal{T}_n\}$ on some time slab \mathcal{T}_n . We define the *multi-adaptive efficiency index* μ , by

$$\mu = \frac{N/k_{\min}}{|\mathcal{T}_n|/k_{\max}} = \frac{k_{\max}}{k_{\min}} \frac{N}{|\mathcal{T}_n|}, \quad (18)$$

where $k_{\min} = \min_{I_{ij} \in \mathcal{T}_n} k_{ij}$, $k_{\max} = \max_{I_{ij} \in \mathcal{T}_n} k_{ij}$, and $|\mathcal{T}_n|$ is the number of local intervals in the time slab \mathcal{T}_n . Thus, to obtain the multi-adaptive efficiency index, we divide the number of local intervals per unit time for a mono-adaptive discretization with the actual number of local intervals per unit time for a multi-adaptive discretization. This is the potential speedup when compared to a mono-adaptive method that is forced to use the same small time step k_{\min} for all components. However, the actual speedup is always smaller than μ for two reasons. The first is the overhead of the multi-adaptive implementation, and the second is that the system of discrete equations on each time slab may sometimes be more expensive to solve than the corresponding mono-adaptive systems (because they are typically larger in size).

Consider a model problem consisting of $N = N_K + N_k$ components, where N_K components vary on a slow time scale K , and N_k components vary on a fast time scale k , as in Figure 4. The potential speedup is given by the multi-adaptive efficiency index,

$$\mu = \frac{K}{k} \frac{N}{N_K + N_k K/k} = \frac{K}{k} \frac{N/K}{N_K/K + N_k/k} \sim \frac{K}{k} \gg 1, \quad (19)$$

if $N_K/K \gg N_k/k$ and $K \gg k$, that is the number of large elements dominates the number of small elements. Thus, the potential speedup can be very large for a system where a large part of the system varies on a large time scale and a small part of the system varies on a small time scale.

If, on the other hand, $K \sim k$ or $N_K \sim N_k$, then the multi-adaptive efficiency index may be of moderate size. As a consequence, the actual speedup may be small (or even negative) if the overhead of the multi-adaptive implementation is significant. In the next section, we indicate the multi-adaptive efficiency index and compare this to the actual speedup for a number of benchmark problems.

6. NUMERICAL EXAMPLES AND BENCHMARK RESULTS

In this section, we present two benchmark problems to demonstrate the efficiency of multi-adaptive time-stepping. Both examples are time-dependent PDEs that we discretize in space using the cG(1) finite element method to obtain a system of ODEs, sometimes referred to as the method of lines approach. In each case, we lump and invert the mass matrix so as to obtain a system of the form (1).

In the first of the two benchmark problems, the individual time steps are chosen automatically, based on an a posteriori error estimate as discussed in Section 3.4. For the second problem, the time steps are fixed in time and determined according to a local CFL condition $k \sim h$, on each element. The results were obtained with DOLFIN version 0.6.2.

6.1 A Nonlinear Reaction-Diffusion Equation

As a first example, we solve the following nonlinear reaction-diffusion equation, taken from Savcenco et al. [2005]:

$$\begin{aligned} u_t - \epsilon u_{xx} &= \gamma u^2(1 - u) && \text{in } \Omega \times (0, T], \\ \partial_n u &= 0 && \text{on } \partial\Omega \times (0, T], \\ u(\cdot, 0) &= u_0 && \text{in } \Omega, \end{aligned} \tag{20}$$

with $\Omega = (0, L)$, $\epsilon = 0.01$, $\gamma = 1000$ and final time $T = 1$.

The equation is discretized in space with the standard cG(1) method using a uniform mesh with 1000 mesh points. The initial data is chosen according to:

$$u_0(x) = \frac{1}{1 + \exp(\lambda(x - 1))}. \tag{21}$$

The resulting solution is a reaction front, sweeping across the domain from left to right, as demonstrated in Figure 5. The multi-adaptive time steps are automatically selected to be small in and around the reaction front, and sweep the domain at the same velocity as the reaction front, as demonstrated in Figure 6.

To study the performance of the multi-adaptive solver, we compute the solution for a range of tolerances with $L = 5$ and compare the resulting error and CPU time with a standard mono-adaptive solver that uses equal adaptive time steps for all components. To make the comparison fair, we compare the multi-adaptive mcG(q) method with the mono-adaptive cG(q) method. In the benchmarks, we only examine $q = 1$. Both methods are implemented for general order q in the same programming language (C++), within a common framework (DOLFIN), but the mono-adaptive method takes full advantage of the fact that the time steps are equal for all components. In particular, the mono-adaptive solver may use much simpler data structures (a plain C array) to store the solution on each time slab and there is no overhead for interpolation of the solution. Furthermore, for the multi-adaptive solver, we need to supply a right-hand side function f which may be called to evaluate single components $f_i(U(t), t)$, while for the mono-adaptive solver, we may evaluate all components of f at the same time, which is usually an advantage for the mono-adaptive solver.

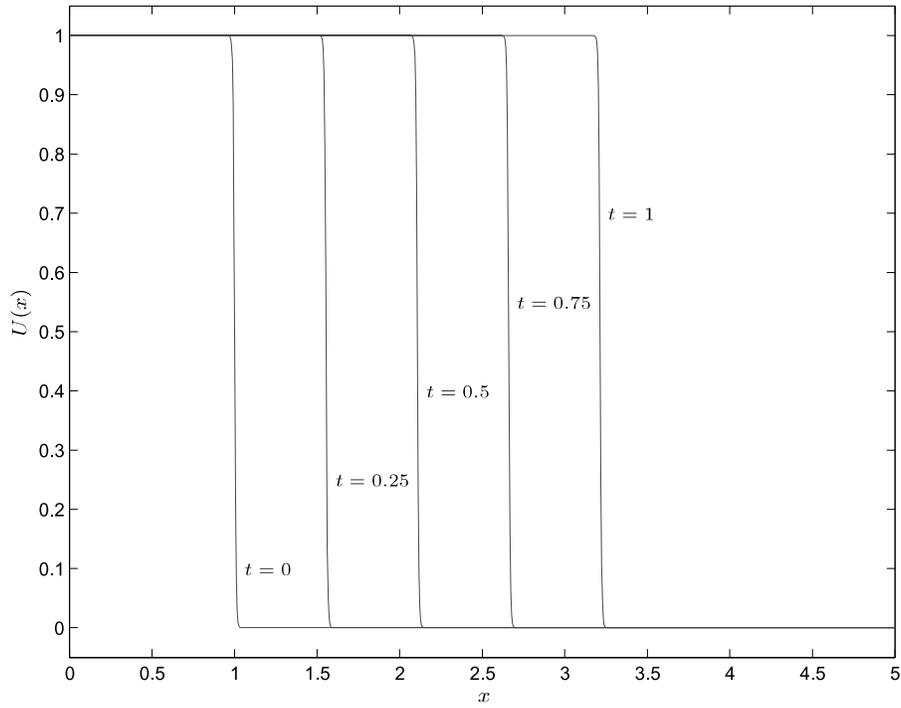


Fig. 5. Propagation of the solution of the reaction–diffusion problem (20).

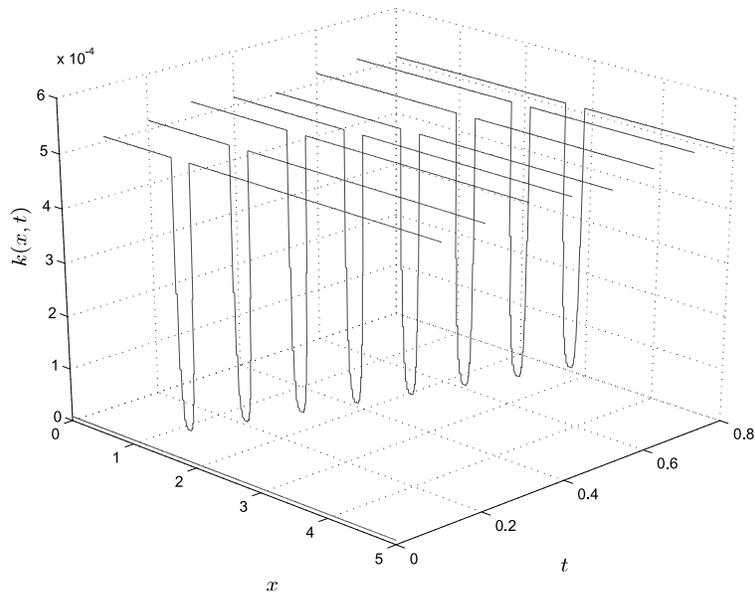


Fig. 6. The multi-adaptive time steps as function of space at a sequence of points in time for the test problem (20).

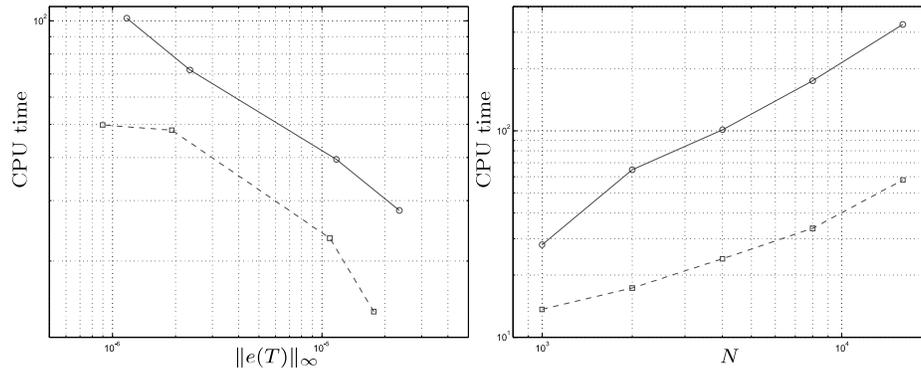


Fig. 7. CPU time as function of the error (left) and number of components N (right) for mcG(1) (dashed line) and cG(1) (solid line) for the test problem (20).

This is a more meaningful measure of performance compared to only measuring the number degrees of freedom (local steps) or comparing the CPU time against the same multi-adaptive solver when it is forced to use identical time steps for all components, as in Logg [2003a], since one must also take into account the overhead of the more complicated algorithms and data structures necessary for the implementation of multi-adaptive time-stepping.

Note that we do not solve the dual problem for computing stability factors (or stability weights), which is necessary to obtain a reliable error estimate. Thus, the tolerance only controls the size of the error modulo the stability factor, which is unknown.

In addition, we also compare the two methods for varying size L , of the domain Ω , keeping the same initial conditions but scaling the number of mesh points according to the length of the domain, $N = 1000L/5$. As the size of the domain increases, we expect the relative efficiency of the multi-adaptive method to increase, since the number of inactive components increases relative to the number of components located within the reaction front.

In Figure 7, we plot the CPU time as function of the tolerance and number of components (size of domain) for the mcG(1) and cG(1) methods. We also summarize the results in Table II and Table III. As expected, the speedup expressed as the multi-adaptive efficiency index μ , that is, the ideal speedup if the cost per degree of freedom were the same for the multi- and mono-adaptive methods, is large in all test cases: about a factor 100. The speedup in terms of the total number of time slabs is also large. Note that in Table II, the total number of time slabs M , remains practically constant as the tolerance and the error are decreased. The decreased tolerance instead results in finer local resolution of the reaction front, which is evident from the increasing multi-adaptive efficiency index. At the same time, the mono-adaptive method needs to decrease the time step for all components, so the relative efficiency of the multi-adaptive method increases as the tolerance decreases. See also Figure 8 for a comparison of the multi-adaptive time steps at two different tolerances.

Table II. Benchmark Results for mcG(1) (top) and cG(1) (bottom) for Varying Tolerance and Fixed Number of Components $N = 1000$ for the Test Problem (20). The Table Shows the Tolerance TOL Used for the Computation, the Error $\|e(T)\|_\infty$ in the Maximum Norm at the Final Time, the Time Used to Compute the Solution, the Number of Time Slabs M (with the Number of Rejected Time Slabs in Parenthesis), the Average Number of Iterations n on the Time Slab System (with the Number of Local Iterations on Sub-Slabs in Parenthesis), and the Multi-Adaptive Efficiency Index, μ

TOL	$\ e(T)\ _\infty$	CPU time	M	n	μ
$1.0 \cdot 10^{-6}$	$1.8 \cdot 10^{-5}$	14.2 s	1922 (5)	3.990 (1.498)	95.3
$5.0 \cdot 10^{-7}$	$1.1 \cdot 10^{-5}$	23.3 s	1912 (9)	4.822 (1.544)	138.2
$1.0 \cdot 10^{-7}$	$1.9 \cdot 10^{-6}$	48.1 s	1929 (7)	4.905 (1.594)	142.6
$5.0 \cdot 10^{-8}$	$9.0 \cdot 10^{-7}$	49.8 s	1917 (7)	4.131 (1.680)	172.4

TOL	$\ e(T)\ _\infty$	time	M	n	μ
$1 \cdot 10^{-6}$	$2.3 \cdot 10^{-5}$	28.1 s	117089 (1)	4.0	1.0
$5 \cdot 10^{-7}$	$1.2 \cdot 10^{-5}$	39.5 s	165586 (1)	4.0	1.0
$1 \cdot 10^{-7}$	$2.3 \cdot 10^{-6}$	71.9 s	370254 (1)	3.0	1.0
$5 \cdot 10^{-8}$	$1.2 \cdot 10^{-6}$	101.7 s	523615 (1)	3.0	1.0

Table III. Benchmark Results for mcG(1) (top) and cG(1) (bottom) for Fixed Tolerance TOL = $1.0 \cdot 10^{-6}$ and Varying Number of Components (and Size of Domain). (See Table II for an Explanation of Table Legends)

N	$\ e(T)\ _\infty$	CPU time	M	n	μ
1000	$1.8 \cdot 10^{-5}$	13.6 s	1922 (5)	4.0 (1.5)	95.3
2000	$1.7 \cdot 10^{-5}$	17.3 s	1923 (5)	4.0 (1.2)	140.5
4000	$1.6 \cdot 10^{-5}$	24.0 s	1920 (6)	4.0 (1.0)	185.0
8000	$1.7 \cdot 10^{-5}$	33.7 s	1918 (5)	4.0 (1.0)	218.8
16000	$1.7 \cdot 10^{-5}$	57.9 s	1919 (5)	4.0 (1.0)	240.0

N	$\ e(T)\ _\infty$	time	M	n	μ
1000	$2.3 \cdot 10^{-5}$	28.1 s	117089 (1)	4.0	1.0
2000	$2.2 \cdot 10^{-5}$	64.8 s	117091 (1)	4.0	1.0
4000	$2.2 \cdot 10^{-5}$	101.3 s	117090 (1)	4.0	1.0
8000	$2.2 \cdot 10^{-5}$	175.1 s	117089 (1)	4.0	1.0
16000	$2.2 \cdot 10^{-5}$	327.7 s	117089 (1)	4.0	1.0

The situation is slightly different in Table III, where the tolerance is kept constant, but the size of the domain and number of components vary. Here, the number of time slabs remains practically constant for both methods, but the multi-adaptive efficiency index increases as the size of the domain increases, since the reaction front then becomes more and more localized relative to the size of the domain. As a result, the efficiency index of the multi-adaptive method increases as the size of the domain is increased.

In all test cases, the multi-adaptive method is more efficient than the standard mono-adaptive method even when the CPU time (wall-clock time) is chosen as a metric for the comparison. In the first set of test cases, with varying tolerance, the actual speedup is about a factor 2.0 whereas in the second test case, with varying size of the domain, the speedup increases from about a factor 2.0 to a factor 5.7 for the range of test cases. These are significant speedups, although far from the ideal speedup which is given by the multi-adaptive efficiency index.

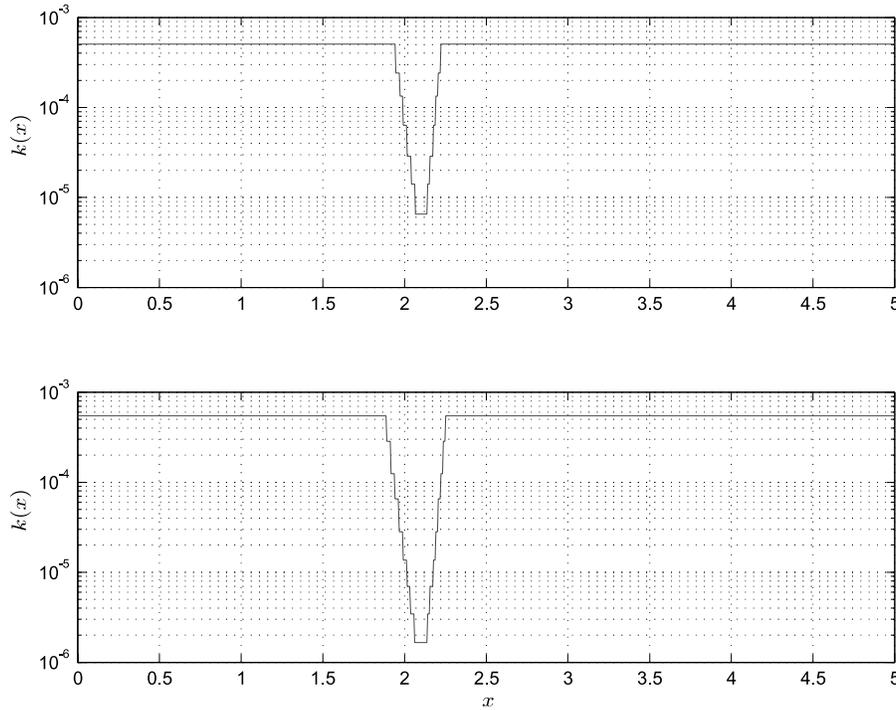


Fig. 8. Multi-adaptive time steps at $t = 0.5$ for two different tolerances for the test problem (20).

There are two main reasons that make it difficult to attain full speedup. The first is that as the size of the time slab increases, the number of iterations n , needed to solve the system of discrete equations increases. In Table III, the number of iterations, including local iterations on individual elements as part of a global iteration on the time slab, is about a factor 1.5 larger for the multi-adaptive method. However, the main overhead lies in the more straightforward implementation of the mono-adaptive method compared to the more complicated data structures needed to store and interpolate the multi-adaptive solution. For constant time step and equal time step for all components, this overhead is roughly a factor 5 for the test problem, but the overhead increases to about a factor 100 when the time slab is locally refined. It thus remains important to further reduce the overhead of the implementation in order to increase the range of problems where the multi-adaptive methods give a positive speedup.

6.2 The Wave Equation

Next, we consider the wave equation,

$$\begin{aligned}
 u_{tt} - \Delta u &= 0 && \text{in } \Omega \times (0, T], \\
 \partial_n u &= 0 && \text{on } \partial\Omega \times (0, T], \\
 u(\cdot, 0) &= u_0 && \text{in } \Omega,
 \end{aligned} \tag{22}$$

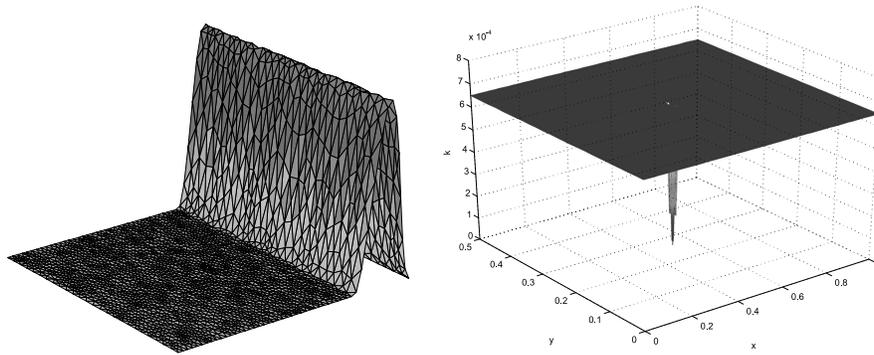


Fig. 9. Initial data (left) and multi-adaptive time steps (right) for the solution of the wave equation.

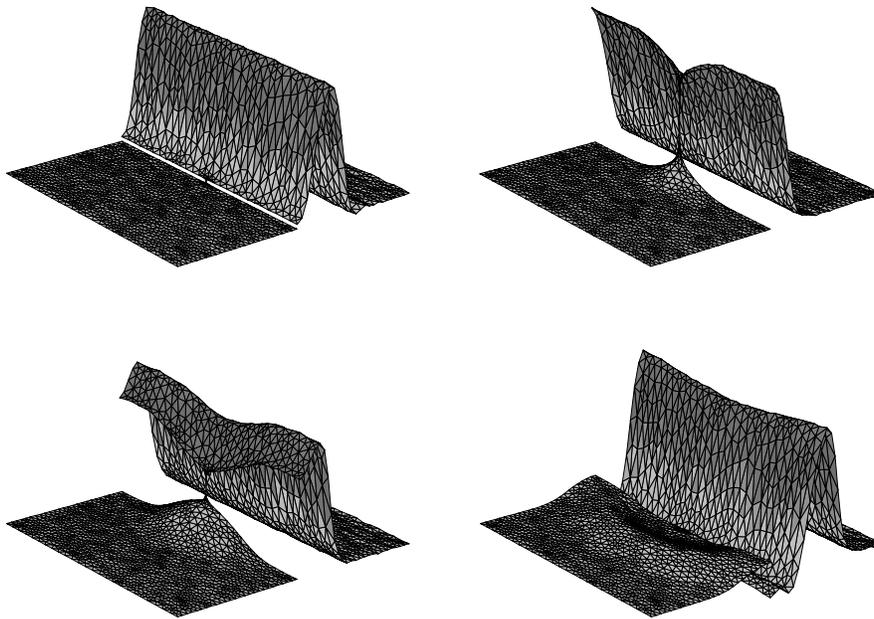


Fig. 10. The solution of the wave equation at times $t = 0.25$, $t = 0.4$, $t = 0.45$, and $t = 0.6$.

on a two-dimensional domain Ω , consisting of two square sub-domains of side length 0.5 separated by a thin wall with a narrow slit of size 0.0001×0.0001 at its center. The initial condition is chosen as a plane wave traversing the domain from right to left. In Figure 9, we plot the initial data together with the fixed multi-adaptive time steps. The resulting solution is shown in Figure 10.

The geometry of the domain Ω forces the discretization to be very fine close to the narrow slit. Further away from the slit, we let the mesh be coarse. The mesh was created by specifying a mesh size h with $h \gg w$ where w is the width of the narrow slit (see Figure 11). We note that for the multi-adaptive efficiency index μ defined in (19) to be large, the total number of elements must

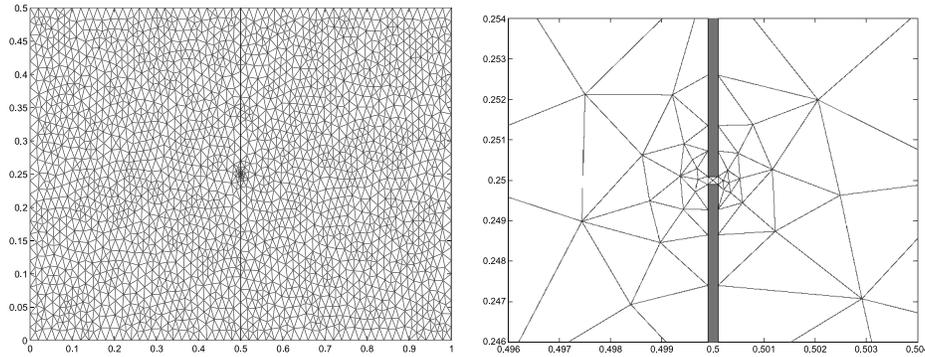


Fig. 11. The mesh used for the solution of the wave equation on a domain intersected by a thin wall with a narrow slit (left) and details of the mesh close to the slit (right).

be large in comparison to the number of small elements close to the narrow slit. Furthermore, the average mesh size must be large compared to the mesh size close to the narrow slit.

For a mono-adaptive method, a global CFL condition puts a limit on the size of the global time step, roughly given by:

$$k \leq h_{\min} = \min_{x \in \Omega} h(x), \quad (23)$$

where $h = h(x)$ is the local mesh size. With a larger time step, an explicit method will be unstable or, correspondingly, direct fixed-point iteration on the system of discrete equations on each time slab will not converge without suitable stabilization.

On the other hand, with a multi-adaptive method, the time step may be chosen to satisfy the CFL condition only locally, that is,

$$k(x) \leq h(x), \quad x \in \Omega, \quad (24)$$

and as a result, the number of local steps may decrease significantly, depending on the properties of the mesh. In this case, with $k = 0.1h$, the speedup for the multi-adaptive mcG(1) method was a factor of 4.2.

7. CONCLUSIONS

We have presented algorithms and data structures for multi-adaptive time-stepping, including the recursive construction of time slabs and efficient interpolation of multi-adaptive solutions. The efficiency of the multi-adaptive methods was demonstrated for a pair of benchmark problems. The multi-adaptive methods mcG(q) and mdG(q) are available as components of DOLFIN, together with implementations of the standard mono-adaptive cG(q) and dG(q) methods. The ODE solvers of DOLFIN are currently being integrated with other components of the FEniCS project, in particular the FEniCS Form Compiler (FFC), [Kirby and Logg 2006, 2007], in order to provide reliable, efficient, and automatic integration of time dependent PDEs.

REFERENCES

- ALEXANDER, S. G. AND AGNOR, C. B. 1998. n -body simulations of late stage planetary formation with a simple fragmentation model. *ICARUS* 132, 113–124.
- DAVÉ, R., DUBINSKI, J., AND HERNQUIST, L. 1997. Parallel treeSPH. *New Astronomy* 2, 277–297.
- DAWSON, C. AND KIRBY, R. C. 2001. High resolution schemes for conservation laws with locally varying time steps. *SIAM J. Sci. Comput.* 22, 6, 2256–2281.
- DELFOUR, M., HAGER, W., AND TROCHU, F. 1981. Discontinuous Galerkin methods for ordinary differential equations. *Math. Comp.* 36, 455–473.
- DUPONT, T., HOFFMAN, J., JOHNSON, C., KIRBY, R. C., LARSON, M. G., LOGG, A., AND SCOTT, L. R. 2003. The FEniCS project. Tech. Rep. 2003–21, Chalmers Finite Element Center Preprint Series.
- ENGSTLER, C. AND LUBICH, C. 1997. MUR8: a multirate extension of the eighth-order Dormand-Prince method. *Appl. Numer. Math.* 25, 2-3, 185–192. Special issue on time integration (Amsterdam, 1996).
- ERIKSSON, K., ESTEP, D., HANSBO, P., AND JOHNSON, C. 1995. Introduction to adaptive methods for differential equations. *Acta Numerica* 4, 105–158.
- ERIKSSON, K. AND JOHNSON, C. 1991. Adaptive finite element methods for parabolic problems I: A linear model problem. *SIAM J. Numer. Anal.* 28, No. 1, 43–77.
- ERIKSSON, K. AND JOHNSON, C. 1995a. Adaptive finite element methods for parabolic problems II: Optimal order error estimates in $l_\infty l_2$ and $l_\infty l_\infty$. *SIAM J. Numer. Anal.* 32, 706–740.
- ERIKSSON, K. AND JOHNSON, C. 1995b. Adaptive finite element methods for parabolic problems IV: Nonlinear problems. *SIAM J. Numer. Anal.* 32, 1729–1749.
- ERIKSSON, K. AND JOHNSON, C. 1995c. Adaptive finite element methods for parabolic problems V: Long-time integration. *SIAM J. Numer. Anal.* 32, 1750–1763.
- ERIKSSON, K., JOHNSON, C., AND LARSSON, S. 1998. Adaptive finite element methods for parabolic problems VI: Analytic semigroups. *SIAM J. Numer. Anal.* 35, 1315–1325.
- ERIKSSON, K., JOHNSON, C., AND THOMÉE, V. 1985. Time discretization of parabolic problems by the discontinuous Galerkin method. *RAIRO MAN* 19, 611–643.
- ESTEP, D. 1995. A posteriori error bounds and global error control for approximations of ordinary differential equations. *SIAM J. Numer. Anal.* 32, 1–48.
- ESTEP, D. AND FRENCH, D. 1994. Global error control for the continuous Galerkin finite element method for ordinary differential equations. *M²AN* 28, 815–852.
- ESTEP, D., LARSON, M., AND WILLIAMS, R. 2000. Estimating the error of numerical solutions of systems of nonlinear reaction–diffusion equations. *Memoirs of the American Mathematical Society* 696, 1–109.
- ESTEP, D. AND STUART, A. 2002. The dynamical behavior of the discontinuous Galerkin method and related difference schemes. *Math. Comp.* 71, 1075–1103.
- ESTEP, D. AND WILLIAMS, R. 1996. Accurate parallel integration of large sparse systems of differential equations. *Math. Models. Meth. Appl. Sci.* 6, 535–568.
- FENICS. 2008. The FEniCS project. www.fenics.org.
- FLAHERTY, J. E., LOY, R. M., SHEPHARD, M. S., SZYMANSKI, B. K., TERESCO, J. D., AND ZIANTZ, L. H. 1997. Adaptive local refinement with octree load balancing for the parallel solution of three-dimensional conservation laws. *J. Para. Distrib. Comput.* 47, 139–152.
- FREE SOFTWARE FOUNDATION. 1999. GNU LGPL. URL: <http://www.gnu.org/licenses/old-licenses/lgpl-2.1.html>.
- GUSTAFSSON, K., LUNDH, M., AND SÖDERLIND, G. 1988. A PI stepsize control for the numerical solution of ordinary differential equations. *BIT* 28, 270–287.
- HOFFMAN, J. AND LOGG, A. 2002. DOLFIN: Dynamic Object oriented Library for FINite element computation. Tech. Rep. 2002–06, Chalmers Finite Element Center Preprint Series.
- HUGHES, T. J. R., LEVIT, I., AND WINGET, J. 1983a. Element-by-element implicit algorithms for heat-conduction. *J. Eng. Mech.-ASCE* 109, 576–585.
- HUGHES, T. J. R., LEVIT, I., AND WINGET, J. 1983b. An element-by-element solution algorithm for problems of structural and solid mechanics. *Comput. Meth. Appl. Mech. Eng.* 36, 241–254.

- HULME, B. L. 1972a. Discrete Galerkin and related one-step methods for ordinary differential equations. *Math. Comput.* 26, 120, 881–891.
- HULME, B. L. 1972b. One-step piecewise polynomial Galerkin methods for initial value problems. *Math. Comput.* 26, 118, 415–426.
- JAMET, P. 1978. Galerkin-type approximations which are discontinuous in time for parabolic equations in a variable domain. *SIAM J. Numer. Anal.* 15, 5, 912–928.
- JOHNSON, C. 1988. Error estimates and adaptive time-step control for a class of one-step methods for stiff ordinary differential equations. *SIAM J. Numer. Anal.* 25, 4, 908–926.
- KIRBY, R. C. AND LOGG, A. 2006. A compiler for variational forms. *ACM Trans. Math. Softw.* 32, 3, 417–444.
- KIRBY, R. C. AND LOGG, A. 2007. Efficient compilation of a class of variational forms. *ACM Trans. Math. Softw.* 33, 3.
- LEW, A., MARSDEN, J. E., ORTIZ, M., AND WEST, M. 2003. Asynchronous variational integrators. *Arch. Rational. Mech. Anal.* 167, 85–146.
- LOGG, A. 2003a. Multi-adaptive Galerkin methods for ODEs I. *SIAM J. Sci. Comput.* 24, 6, 1879–1902.
- LOGG, A. 2003b. Multi-adaptive Galerkin methods for ODEs II: Implementation and applications. *SIAM J. Sci. Comput.* 25, 4, 1119–1141.
- LOGG, A. 2004. Automation of computational mathematical modeling. Ph.D. thesis, Chalmers University of Technology, Sweden.
- LOGG, A. 2006. Multi-adaptive Galerkin methods for ODEs III: A priori error estimates. *SIAM J. Numer. Anal.* 43, 6, 2624–2646.
- LOGG, A. 2007. Automating the finite element method. *Arch. Comput. Methods Eng.* 14, 93–138.
- LOGG, A., WELLS, G. ET AL. DOLFIN: A general-purpose finite element library. <http://www.fenics.org/dolfin/>.
- MAKINO, J. AND AARSETH, S. 1992. On a Hermite integrator with Ahmad-Cohen scheme for gravitational many-body problems. *Publ. Astron. Soc. Japan* 44, 141–151.
- OSHER, S. AND SANDERS, R. 1983. Numerical approximations to nonlinear conservation laws with locally varying time and space grids. *Math. Comp.* 41, 321–336.
- SAVCENCO, V. 2008. Multirate numerical integration for ordinary differential equations. Ph.D. thesis, Universiteit van Amsterdam.
- SAVCENCO, V., HUNSDORFER, W., AND VERWER, J. 2005. A multirate time stepping strategy for parabolic PDEs. Tech. Rep. MAS–E0516, Centrum voor Wiskunde en Informatica (CWI).
- SÖDERLIND, G. 2003. Digital filters in adaptive time-stepping. *ACM Trans. Math. Softw.* 29, 1, 1–26.

Received June 2006; revised February 2008; accepted April 2008