Automatic Differentiation of Numerical Integration Algorithms^{*}

Peter Eberhard Institute B of Mechanics University of Stuttgart 70550 Stuttgart Germany pe@mechb.uni-stuttgart.de Christian Bischof Argonne National Laboratory Mathematics and Computer Science Division Argonne, IL 60439-4844, U.S.A. bischof@mcs.anl.gov

Argonne Preprint ANL/MCS-P621-1196

Abstract

Automatic differentiation (AD) is a technique for automatically augmenting computer programs with statements for the computation of derivatives. This article discusses the application of automatic differentiation to numerical integration algorithms for ordinary differential equations (ODEs), in particular, the ramifications of the fact that AD is applied not only to the solution of such an algorithm, but to the solution procedure itself. This subtle issue can lead to surprising results when AD tools are applied to variable-stepsize, variable-order ODE integrators. The computation of the final time step plays a special role in determining the computed derivatives. We investigate these issues using various integrators and suggest constructive approaches for obtaining the desired derivatives.

1 Introduction

Typically, the description of technical systems or natural phenomena leads to complicated mathematical models involving ordinary differential equations (ODEs), differential-algebraic equations, or partial differential equations.

For example, many problems in mechanical, chemical, and electrical engineering can be formulated as an initial value problem using ODEs:

For given values of system parameters $p \in \mathbb{R}^h$, find the trajectories $x(t, p), x \in \mathbb{R}^n$ for $t^0 \leq t \leq t^1$, where x is the state vector, t

^{*}Mathematics Subject Classification: 34A12, 65L05, 65L06, 68N99

the time, t^0 the initial time, and t^1 the final time. The state is determined by the solution of the initial value problem

$$\mathbf{x} = \mathbf{f}(\mathbf{x}, \mathbf{p}, t), \quad \mathbf{x}(t = t^0, \mathbf{p}) = \mathbf{x}^0, \tag{1}$$

where f is the vector of state derivatives and x^0 is the initial state.

Problem (1) typically is solved by using a numerical integration algorithm, and a large body of literature is devoted to this subject (see, e.g., [6, 13, 14, 15, 19, 20, 21]). Also, in many engineering applications, one is interested not only in the final state, but also in performance criteria ψ computed from the trajectories \boldsymbol{x} .

If optimization procedures are applied in order to choose optimal design variables with respect to certain performance criteria, or if parameter estimation techniques are used to identify model parameters from measurements (see, e.g., [2, 8]), then, with the final state

$$\boldsymbol{x}^1 := \boldsymbol{x}(t^1, \boldsymbol{p}), \tag{2}$$

one typically has to compute

$$\frac{\partial \boldsymbol{x}^1}{\partial \boldsymbol{p}^T}.$$
(3)

While the sensitivities at the final time step are of interest, their numerical value may depend on the whole time history of the system (e.g., when a performance criterion is some integral over \boldsymbol{x} from t^0 to t^1).

During the past decade, so-called automatic differentiation (AD) tools have been developed that make it possible to augment general Fortran or C codes with statements for the computation of derivatives in an automated fashion. AD relies on the fact that every computer program employs only simple operations such as additions or multiplications or intrinsics such as sin(), whose derivatives are known and then computes derivatives for the whole program by judicious composition of these derivatives using the chain rule of differential calculus. In our experiments, we used the ADIFOR¹ tool [3]; the ADIFOR reference, as well as [1, 10, 17], provides some basics on AD. The impact of the associativity of the chain rule on the cost of computing derivatives is discussed in [4, 11], and a collection of currently available AD tools can be found on the World-Wide Web².

While AD has been used successfully in many large-scale applications and inherently computes accurate derivatives, the black-box application of AD can lead to surprising results, because AD will differentiate not only the solution computed by a computer program, but also the algorithm by which the solution is being derived. That is, while AD will compute a derivative, the value of this derivative may well depend on the algorithm chosen to compute the solution.

¹See http://www.mcs.anl.gov/adifor and http://www.cs.rice.edu/~adifor.

²See http://www.mcs.anl.gov/Projects/autodiff/adtools.

In this article we investigate the automatic differentiation of numerical integration algorithms. In the next section, we consider the automatic differentiation of a prototypical integration algorithm and illustrate the impact that error-adaptive schemes can have on the computed derivatives—that is, different integrators can lead to very different values for the computed derivatives. This realization also leads us to suggest two approaches to suppress the impact of the solution algorithm on derivatives, thus resulting in the computation of derivatives that are defined by the nondiscretized solution $\boldsymbol{x}(t, \boldsymbol{p})$ and, in general, are the desired ones.

We illustrate these effects and our remediation methodology in Section 3 on a relatively simple ODE with known explicit solution, using single-step Runge-Kutta integrators with and without adaptive stepsize control, as well as the multistep Shampine-Gordon algorithm. The computation of the final time step has a fundamental effect on the overall derivatives, and this issue is also investigated. In Section 4, we then apply our approach to a complicated problem from multibody dynamics and verify the results using the adjoint variable approach, an approximation-free method to efficiently compute sensitivities for multibody systems. Lastly, we summarize our results.

2 Automatic Differentiation of Prototypical Numerical Integration Algorithms

The numerical integration of ODEs is one of the basic problems in numerical computing, and many research groups are working on the development of reliable and efficient algorithms. These algorithms can be categorized into several classes, including single step algorithms, multistep algorithms, and extrapolation algorithms.

Special algorithms also exist, for example, for stiff systems, highly or loosely coupled systems, and systems composed of several subsystems with different frequency ranges or real-time requirements.

2.1 Derivatives of Time

To illustrate the issues relevant to the interplay of AD and integration algorithms, we choose some explicit single-step algorithms of Euler and Runge-Kutta type with and without stepsize control, as well as a sophisticated multistep integration algorithm with adaptive stepsize and interpolation order control. The following discussion of integration algorithms is intentionally kept simple to emphasize the salient points. Details on numerical integration algorithms and their implementation can be found, for example, in the aforementioned references and in the many codes available from netlib.³

³See http://netlib.att.com/netlib/master/readme.html.

With single-step algorithms, the time discretization that typically is applied to solve (1) leads to a recursive scheme

$$\mathbf{x}_{i+1} = \mathbf{x}_i + h_i \dot{\mathbf{x}}_i, \qquad t_{i+1} = t_i + h_i,$$
(4)

where the subscript *i* denotes the *i*th integration step. That is, $\mathbf{x}_i := \mathbf{x}(t_i)$, h_i is the actual stepsize, and \mathbf{x}_i denotes a slope estimation. The simplest case $h_i = h = constant$ and $\mathbf{x}_i = \mathbf{x}_i$ yields the explicit Euler scheme; with $\mathbf{x}_i = \mathbf{x}_{i+1}$, on the other hand, we have the implicit Euler scheme. Usually \mathbf{x}_i is composed from different evaluations of the ODE at different times and approximations. The (nonunique) weighting coefficients of its different components have to satisfy a Taylor series approximation with certain order.

Most advanced integrators employ adaptive stepsize control that, based on local error estimates (e.g., available from a doublestep technique or the simultaneous evaluation of different-order schemes), dynamically adapts the stepsize. This stepsize control is essential for efficiency, enabling one to choose the stepsize as big as possible, yet at the same time sufficiently small enough to restrict the integration error. Multistep algorithms additionally use information from former steps to predict appropriate stepsizes and slopes. Many sophisticated modifications, such as the use of variable extrapolation order or projections, may further improve the efficiency. As a result, all variables in (4) may depend on the system parameters p, that is,

$$\boldsymbol{x}_{i+1}(\boldsymbol{p}) = \boldsymbol{x}_i(\boldsymbol{p}) + h_i(\boldsymbol{p}) \dot{\boldsymbol{x}}_i(\boldsymbol{p}).$$
(5)

The procedure defined by (5) is the blueprint of a numerical algorithm, which starts from the initial values $\mathbf{x}_0 = \mathbf{x}^0$ and the system parameters \mathbf{p} and computes some final values $\mathbf{x}^1(\mathbf{p})$ via an often extremely large number of intermediate steps. Differentiating (5) with respect to \mathbf{p} , with

$$\nabla \boldsymbol{x} := \frac{d\boldsymbol{x}}{d\boldsymbol{p}^T},\tag{6}$$

we obtain

$$\nabla \boldsymbol{x}_{i+1} = \nabla \boldsymbol{x}_i + h_i \nabla \bar{\boldsymbol{x}}_i + \nabla h_i \bar{\boldsymbol{x}}_i.$$
(7)

Note that depending on how h_i is computed, we may obtain very different values for the total derivative ∇x^1 . If the initial values x^0 are independent from p, then $\nabla x^0 = 0$; otherwise, ∇x^0 has nonzero components.

2.2 Computation of the Desired Derivatives for the State Variables

To obtain the desired derivatives, we can consider two choices, which are illustrated in Figure 1.



Figure 1: Manual transformation versus automatic transformation with AD tools

1. Differentiate the ODE and integrate:

Differentiating (1) with respect to p, we obtain with $dt/dp^T = 0$

$$\frac{d}{d\boldsymbol{p}^{T}}(\boldsymbol{\dot{x}}) = \frac{d}{d\boldsymbol{p}^{T}} \left(\frac{d\boldsymbol{x}}{dt}\right) = \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}^{T}} \frac{d\boldsymbol{x}}{d\boldsymbol{p}^{T}} + \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{p}^{T}}.$$
(8)

Exchanging the order of differentiation, we thus obtain a new ODE for ∇x :

$$\frac{d}{dt}\left(\frac{d\boldsymbol{x}}{d\boldsymbol{p}^{T}}\right) = \frac{d}{dt}\left(\nabla\boldsymbol{x}\right) = \left[\nabla\boldsymbol{x}\right] = \frac{\partial\boldsymbol{f}}{\partial\boldsymbol{x}^{T}}\nabla\boldsymbol{x} + \frac{\partial\boldsymbol{f}}{\partial\boldsymbol{p}^{T}}, \quad \nabla\boldsymbol{x}(t=t^{0}) = \nabla\boldsymbol{x}^{0},$$
(9)

which we can integrate alongside our original solution. AD techniques could, for example, be employed to compute $\partial \boldsymbol{f} / \partial \boldsymbol{x}^T$ and $\partial \boldsymbol{f} / \partial \boldsymbol{p}^T$, but we would not differentiate through the integration algorithm for \boldsymbol{x} .

Given a suitable integrator, this approach will deliver the desired sensitivities ∇x^1 with suitable accuracy, since its behavior is governed jointly by (1) and (9). Up to the chosen tolerance, the actual time discretization will not impact either x^1 or ∇x^1 .

2. Differentiate the integrator for the ODE:

Viewing the integration procedure for solving (1) as a black box that, given values for p produces values for x^1 , we can employ an AD tool to augment both the problem-independent code for the numerical time integration algorithm and the problem-specific code for the evaluation of the ODE x = f(x, p, t) with statements for the computation of derivatives.

We observe two facts:

(a) Since AD tools do not alter the control flow of the original program, the time discretization chosen is determined solely by the integration of (1).

(b) As observed in (5), the stepsize h_i is likely to depend on the parameters \boldsymbol{p} , and an AD tool will associate a gradient ∇h_i with h_i . Thus, the update (7), which will be computed by an AD tool, leads to an inconsistent integrator for $\nabla \boldsymbol{x}$, as the final result depends on the discretization strategy chosen. Hence, it is also unlikely that $\nabla \boldsymbol{x}^1 = \nabla \boldsymbol{x} |_{t_i - t_i}$ equals the desired $\partial \boldsymbol{x}^1 / \partial \boldsymbol{p}^T$.

The automatically differentiated integration algorithm computes $x^1(t^1(p), p)$, where the physically implausible dependence of t^1 on p results only from the adaptive time discretization. Differentiating with respect to p, we obtain

$$\nabla \boldsymbol{x}^{1} = \frac{\partial \boldsymbol{x}^{1}}{\partial t^{1}} \nabla t^{1} + \frac{\partial \boldsymbol{x}^{1}}{\partial \boldsymbol{p}^{T}}.$$
(10)

The total derivatives $\nabla \boldsymbol{x}^1 = d\boldsymbol{x}^1/d\boldsymbol{p}^T$ and $\nabla t^1 = dt^1/d\boldsymbol{p}^T$ depend on the time discretization and are computed by the AD-generated code. Thus, to arrive at the desired solution $\partial \boldsymbol{x}^1/\partial \boldsymbol{p}^T$, we can pursue one of the following strategies:

• Perform an a posteriori error correction:

From (10), taking into account (1), we realize that the desired derivatives $\partial x^1 / \partial p^T$, which do not depend on the time discretization chosen, can be computed as

$$\frac{\partial \boldsymbol{x}^{1}}{\partial \boldsymbol{p}^{T}} = \nabla \boldsymbol{x}^{1} - \boldsymbol{f}(\boldsymbol{x}^{1}, \boldsymbol{p}, t^{1}) \nabla t^{1}.$$
(11)

• Use an integrator with fixed stepsize:

In this case, we have $\nabla h_i \equiv 0, \forall i$. Thus $\nabla t \equiv 0$, and hence $\nabla x^1 \equiv \partial x^1 / \partial p^T$ in (10). The AD-computed derivative trajectories are the desired ones, and thus, no modification is required for fixed-stepsize integration algorithms (see also [18], which explores this issue in more detail in the context of a so-called quasi-steady state integrator).

• Modify the AD-generated code to enforce $\nabla h_i = 0, \forall i$:

For the first step, the user must guess an initial stepsize h_0 . Because h_0 is independent of \boldsymbol{p} , we have $\nabla h_0 = 0$. Assume that the correct stepsize is known in advance for each step. Then $\nabla h_i = 0, \forall i$ (which implies $\nabla t_i = 0, \forall i$) and one gets the same correct results as for the fixed stepsize algorithms. Thus, by modifying the AD-generated code manually to ensure

$$\nabla h_i = 0 \tag{12}$$

we can ensure consistency. While this procedure could be done easily when the code is developed, and may in fact be desirable because of the potentially unpredictable nature of $\nabla \boldsymbol{x}$ and ∇t even when $\partial \boldsymbol{x}/\partial \boldsymbol{p}^T$ is well behaved, it is likely to be a nontrivial task after the fact, since in-depth knowledge of the differentiated code may be required in order not to miss subtle dependencies. Perhaps surprising, a commonly used strategy for determining \boldsymbol{x}^1 for a predetermined t^1 results in $\nabla t^1 = 0$ and hence $\frac{\partial \boldsymbol{x}^1}{\partial \boldsymbol{p}^T} = \nabla \boldsymbol{x}^1$. If, for a suggested stepsize h_i , we have in the last step $t_i + h_i > t^1$, we are likely to set

$$h_{i+1} = t^1 - t_i \quad (\text{implying } \nabla h_{i+1} = -\nabla t_i), \tag{13}$$

since t^1 is a user-selected constant. Thus, the computation

$$t_{i+1} = t_i + h_{i+1}$$
 implies $\nabla t^1 = \nabla t_{i+1} = \nabla t_i + \nabla h_{i+1} = 0.$ (14)

Hence, for the last step (and most likely only for the last step), we have $\nabla t_i = 0$, and therefore the a posteriori correction is not required for ∇x^1 , although it is most likely required for any other point on the derivative trajectory.

However, there is no guarantee that $\nabla t^1 = 0$ if t^1 has been preselected. Another frequently chosen approach to compute \boldsymbol{x}^1 is to terminate the integration at the first time point beyond t^1 and then interpolate the value of \boldsymbol{x} for t^1 . In this case, ∇t^1 is unlikely to be zero, and the error correction is required. Thus, for the general user, who most likely is not familiar with the internals of the algorithm, we suggest checking whether $\nabla t|_{t=t^1}$ is zero. If it is not, the a posteriori correction (11) should be applied.

3 Experimental Results with a One-Mass Oscillator

The simplest multibody system is a horizontal one-mass oscillator shown in Figure 2. As shown, for example, in [2], one can derive closed-form solutions for both the state and its gradients, and thus the system is well suited as an example to illustrate the issues outlined in the preceding section.



Figure 2: One-mass oscillator

3.1 Mechanical Model

A body with mass m can slide on the horizontal ground. It is coupled to the wall with a linear spring with spring stiffness c. The position is described by

y(t). From Newton's equation one can derive the equation of motion

$$m\ddot{y} + cy = 0 \tag{15}$$

or, with $\boldsymbol{x} = [y, \dot{y}]^T$, the corresponding set of first-order ODEs

$$\dot{\boldsymbol{x}} = \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -\frac{c}{m} x_1 \end{bmatrix}.$$
(16)

With the initial condition $\boldsymbol{x}(t=t^0) = \boldsymbol{x}^0 = [0, v^0]^T$, the solution of the ODE is

$$y(t) = v^0 \sqrt{\frac{m}{c}} \sin \sqrt{\frac{c}{m}} t.$$
(17)

For m = 1 and the system parameters $\boldsymbol{p} = [c, v^0]^T$, we find

$$y(t) = v^0 \frac{1}{\sqrt{c}} \sin \sqrt{c}t, \quad \dot{y}(t) = v^0 \cos \sqrt{c}t, \quad \ddot{y}(t) = -v^0 \sqrt{c} \sin \sqrt{c}t.$$
 (18)

Two criteria are now defined: The criterion ψ_1 contains the position of the body at time $t^1 = \pi/2$, and for $\boldsymbol{p} = [1, 0.5]^T$ we have

$$\frac{d\psi_1}{dp_1} = \frac{d\psi_1}{dc} = \frac{v^0}{2c} \left(\frac{\pi}{2} \cos\left(\sqrt{c}\frac{\pi}{2}\right) - \frac{1}{\sqrt{c}} \sin\left(\sqrt{c}\frac{\pi}{2}\right) \right) = -0.25,$$

$$\frac{d\psi_1}{dp_2} = \frac{d\psi_1}{dv^0} = \frac{1}{\sqrt{c}} \sin\left(\sqrt{c}\frac{\pi}{2}\right) = 1.0.$$
 (19)

The criterion ψ_2 integrates the position over the whole interesting simulation time interval $[t^0, t^1]$

$$\psi_2 = \int_{t^0}^{t^1} y(t) \, dt = \int_{t^0=0}^{t^1=\pi/2} \frac{v^0}{\sqrt{c}} \sin\left(\sqrt{c}t\right) dt = \frac{v^0}{\sqrt{c}} \left(1 - \cos\sqrt{c}\frac{\pi}{2}\right). \tag{20}$$

Here, the integral type criterion can still be computed analytically, but for more complicated systems it has to be evaluated numerically together with the ODE of the mechanical system, yielding an extended state

$$\mathbf{x} = \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{bmatrix} = \begin{bmatrix} x_2 \\ -\frac{c}{m} x_1 \\ \dot{\psi}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -\frac{c}{m} x_1 \\ x_1 \end{bmatrix}.$$
 (21)

Again we find explicit solutions for the gradients

$$\frac{d\psi_2}{dp_1} = \frac{d\psi_2}{dc} = \frac{\frac{v^0}{2\sqrt{c}}\frac{\pi}{2}\sin\left(\sqrt{c}\frac{\pi}{2}\right) + v^0\cos\left(\sqrt{c}\frac{\pi}{2}\right)}{c^2} - \frac{v^0}{c^2},$$
(22)

$$\frac{d\psi_2}{dp_2} = \frac{d\psi_2}{dv^0} = \frac{1}{c} - \frac{1}{c} \cos\left(\sqrt{c}\frac{\pi}{2}\right),$$
(23)

and, for the given numerical values, the results

$$\frac{d\psi_2}{dp_1} = \frac{\pi}{8} - \frac{1}{2}, \qquad \frac{d\psi_2}{dp_2} = 1.0.$$
(24)

3.2 Single-Step Integration Algorithms without Stepsize Control

We investigate three similar integration schemes: the explicit Euler scheme, the Heun algorithm, and the fourth-order Runge-Kutta algorithm; see [16]. Because the stepsize h does not change during the integration, ADIFOR does not generate a gradient ∇h .

Table 1 shows the relative errors for the criteria and the gradients for different stepsizes and integration algorithms. Only minor differences between the reference gradient and the AD gradients exist. The relative error in the criteria is about the size of the error in the gradients. As expected, automatic differentiation of single-step integration algorithms without stepsize control leads to the desired result without any need for user modifications.

The differences can be explained by the choice of the stepsize and the algorithm; no additional errors are introduced by the automatic differentiation procedure. The higher-order algorithms yield higher accuracies, and the errors are at least of order O(h) for the Euler integrator, $O(h^2)$ for the Heun integrator and $O(h^4)$ for the Runge-Kutta integrator. Some components yield even higher error order, for example, for the Runge-Kutta integrator ψ_1 and $d\psi_1/dp_2$ seems to be of order $O(h^5)$.

Step h	Method	ψ_1	ψ_2	$\frac{d\psi_1}{dp_1}$	$\frac{d\psi_1}{dp_2}$	$\frac{d\psi_2}{dp_1}$	$\frac{d\psi_2}{dp_2}$
0.25	Euler	.20e-0	.36e-1	.28e-0	.20e-0	.63e-0	.36e-1
	Heun	.28e-2	.15e-1	.16e-1	.28e-2	.45e-1	.15e-1
	Ru-Ku	.10e-4	.48e-4	.25e-4	.10e-4	.37e-3	.48e-4
0.025	Euler	.20e-1	.33e-3	.21e-1	.20e-1	.72e-1	.33e-3
	Heun	.30e-5	.16e-3	.25e-3	.30e-5	.40e-3	.16e-3
	Ru-Ku	.11e-9	.51e-8	.74e-8	.11e-9	.36e-7	.51e-8
0.0025	Euler	.20e-2	.33e-5	.20e-2	.20e-2	.72e-2	.33e-5
	Heun	.31e-8	.16e-5	.26e-5	.31e-8	.38e-5	.16e-5
	Ru-Ku	.14e-14	.49e-12	.76e-12	.14e-14	.36e-11	.49e-12

Table 1: Relative errors for different stepsizes

3.3 Single-Step Integration Algorithms with Adaptive Stepsize Control

The next algorithm we investigate is a mixed fourth-/fifth-order Runge-Kutta algorithm with stepsize control; see [7]. To get an estimate for the absolute local error, a fifth-order Runge-Kutta method

$$\boldsymbol{x}_{i+1} = \boldsymbol{x}_i + h(\ldots) + O(h^6) \tag{25}$$

and an embedded fourth-order Runge-Kutta method

$$\boldsymbol{x}_{i+1}^{\star} = \boldsymbol{x}_i + h(\ldots) + O(h^5) \tag{26}$$

are evaluated. The error is of magnitude $O(h^5)$ and follows from the difference

$$\Delta = ||\boldsymbol{x}_{i+1} - \boldsymbol{x}_{i+1}^{\star}||_{\infty}.$$
(27)

Because the error Δ is of magnitude h^5 , we can estimate the required stepsize \bar{h} from the desired error bound $\bar{\Delta}$, the actual stepsize h, and the actual error Δ :

$$\frac{\bar{h}^5}{\bar{\Delta}} \approx \frac{h^5}{\Delta} \quad \to \quad \bar{h} \approx h \sqrt[5]{\frac{\bar{\Delta}}{\Delta}}.$$
(28)

If h > h, the actual stepsize was too big, and the step has to be repeated with decreased stepsize \bar{h} until the error estimate is acceptable. Otherwise, the next (increased) stepsize is computed, and the integration proceeds. Because the actual stepsize h_i and the actual time t_i depend on the state x_i and therefore on the system parameters p, an automatic differentiation tool will compute gradients $\nabla h_i = dh_i/dp^T$ and $\nabla t_i = dt_i/dp^T$, respectively, as suggested in (7).

We then employ relation (11) to compute, at every time step, the desired $\partial \boldsymbol{x}/\partial \boldsymbol{p}^T$ from $\nabla \boldsymbol{x}$ and ∇t . Figure 3 shows some of the trajectories from $\partial \boldsymbol{x}/\partial \boldsymbol{p}^T$, $\nabla \boldsymbol{x}$, and ∇t , where the error tolerance was chosen to be 10^{-8} . Note that the trajectories for $\partial \boldsymbol{x}/\partial \boldsymbol{p}^T$ and the AD-computed $\nabla \boldsymbol{x}$ are very different, but the gradients computed with the correction (11) lead to the correct results. (The derivatives with respect to p_2 and the derivatives related to x_2 and x_3 show the same behavior but, for clarity of presentation, are not drawn here.)

The error can be controlled by the user-defined error bound; see Table 2, where the final time is set to $t^1 = 10$. It can be seen that the relative errors in $\psi_1 = x_1, \psi_2 = x_3$ and their derivatives are both within the prescribed error bound. Of course, the stricter the prescribed error bounds, the higher the number of steps required. In Figure 4 the trajectories are shown for two different prescribed errors bounds $(10^{-5} \text{ and } 10^{-8})$. The trajectories for \dot{x} and $\partial x/\partial p^T$ nearly coincide for different error bounds, but because the time discretization depends on the prescribed error bound, we get significantly different trajectories for ∇x and ∇t .

Table 2: Relative errors in the states for different prescribed error bounds

$\bar{\Delta}$	# Steps	ψ_1	ψ_2	$\frac{d\psi_1}{dp_1}$	$\frac{d\psi_1}{dp_2}$	$\frac{d\psi_2}{dp_1}$	$\frac{d\psi_2}{dp_2}$
10^{-5}	21	.11e-4	.36e-4	.82e-4	.32e-4	.73e-4	.11e-4
10^{-8}	72	.89e-7	.36e-7	.11e-6	.54e-7	$.35\mathrm{e} ext{-}7$.22e-7
10^{-11}	281	.83e-10	.37e-10	.88e-10	.51e-10	.54e-10	.19e-11



Figure 3: Trajectories $\partial \psi_1 / \partial p_1 = \partial x_1 / \partial p_1$, $\nabla \psi_1 = \nabla x_1$, ∇t_1 and $\psi_1 = x_1$

Note that in the unmodified AD-created code the stepsize is controlled only by the integration of the *n* ODEs for \boldsymbol{x}_i and is *not* affected by the *nh* ODEs for $\nabla \boldsymbol{x}_i$. Therefore, the prescribed error bound $\bar{\Delta}$ is valid only for \boldsymbol{x} .

By a manual modification of the generated code, we can also include the extended state ∇x_i in the error estimation. This allows us to guarantee correct results within the error bounds also for the sensitivities, but of course the integration may require more time steps because of the larger dimension n + nh of the extended ODE.

Different methods to compute the criteria and their derivatives at the final timestep have been investigated. The simplest approach is to restrict the size of the final timestep. This can easily be done by checking the new proposed time $t_{i+1} = t_i + h_{i+1}$ after every step. If $t_{i+1} > t^1$, we restrict h_{i+1} to be $t^1 - t_i$. As described in (14), ∇t^1 will be zero and no correction (11) is required to obtain the correct results.

As an alternative, we used interpolation. Here the integration is stopped as



Figure 4: Trajectories for different time discretizations

soon as $t_i > t^1$, and the criteria and derivatives at the final time t^1 are computed by an interpolation between the last two integration steps, for example, the linear interpolation

$$\boldsymbol{x}^{1} = \boldsymbol{x}_{i-1} + (\boldsymbol{x}_{i} - \boldsymbol{x}_{i-1}) \frac{t^{1} - t_{i-1}}{t_{i} - t_{i-1}}.$$
(29)

As expected, the results for the criteria are correct, but the results for their gradients are wrong if the correction (11) is not applied. The use of higherorder interpolation schemes leads to different values for $\nabla \boldsymbol{x}^1$ and ∇t^1 , but the value for $\frac{\partial \boldsymbol{x}^1}{\partial \boldsymbol{p}^T}$ computed from 11 does not change appreciably. If the final time is not known explicitly, but is determined implicitly by a

If the final time is not known explicitly, but is determined implicitly by a final condition $H^1(t^1, \boldsymbol{x}^1, \boldsymbol{p}) = 0$, a root search has to be performed as described in [20]. The final timestep is not actually performed until the root search determined the final time with sufficient accuracy. Therefore, not the root search algorithm itself but only its result influences the final state. Because the final timestep will be chosen to reach t^1 , the results for both the criteria and their gradients need no correction (11). This issue is described in detail in [9].

3.4 Multistep Integration Algorithm

The four algorithms described so far are robust, but for production codes, more sophisticated algorithms typically are used. In multibody dynamics the Shampine-Gordon algorithm [20] serves as standard solver and has proved its reliability and efficiency in many applications. It is a multistep algorithm, in which the information already available from previous steps is used to predict further steps. Not only is the stepsize adjusted adaptively, but also the order of extrapolation polynomials is controlled by local error estimates. For trajectories that are not too rough (i.e., nonstiff systems), high polynomial orders and large stepsizes are obtained.

The integration algorithm consists of about 900 lines of rather complicated Fortran code. Therefore, a manual modification of the code to ensure $\nabla h_i = 0$ is not deemed to be a reliable approach, and the a posteriori error correction (11) is applied to the ADIFOR-generated derivative code. This leads to the correct results; see Table 3.

The correct final time in the investigated version of the Shampine-Gordon algorithm is computed by using a clever interval bisection routine as described in [20]. This guarantees that the evaluation of the last accepted step is at the final time and $\nabla t^1 = 0$.

Error Bound	ψ_1	ψ_2	$\frac{\partial \psi_1}{\partial p_1}$	$\frac{\partial \psi_1}{\partial p_2}$	$\frac{\partial \psi_2}{\partial p_1}$	$\frac{\partial \psi_2}{\partial p_2}$
.1e-2	.11e-2	.66e-2	.12e-1	.11	.11e-1	.18e-1
.1e-5	.17e-4	.66e-5	.16e-5	.82e-4	.22e-4	.23e-4
.1e-8	.15e-7	.15e-8	.59e-8	.79e-7	.14e-7	.60e-8

Table 3: Relative errors in the states for different prescribed error bounds

4 Application to a Technical System

To allow comparisons with analytical results, we kept the example simple, but it was already possible to show the properties of the differentiated integration algorithms and various pitfalls that have to be considered. A more complicated example from robotics will be presented in the following to show that the presented effects also allow a correct handling of interesting real-world problems.

The robot in Figure 5 consists of seven bodies, has five degrees of freedom, (i.e., it allows five independent motions), and is described by ten ODEs. It is described in detail in [2], where the sensitivity of the position of the end effector at the final time with respect to disturbances in several system parameters p =

 $[F_{1z}, L, t_{end}, m_2, I_{3zz}]^T$ is investigated. F_{1z} is a driving force, L a geometrical length, t_{end} the final time, m_2 a mass, and I_{3zz} a component of the inertia tensor. For optimization purposes, additional criteria such as minimal energy consumption or minimal process time are interesting, but for clarity here we restrict ourselves to only one criterion. The results are verified by using the adjoint variable method (AVM) and very costly finite-difference approximations with adaptive order control [2].



Figure 5: Robot

The reference criterion and reference gradient obtained by using the adjoint variable method with integration error tolerances near machine accuracy is for the component $\partial \psi / \partial p_1$ as follows:

$$\psi = -4.136636, \quad \frac{\partial \psi}{\partial p_1} = 0.0186126.$$
(30)

Usually it is not required to compute the gradients to such high accuracy. If the relative and absolute error bounds for the Shampine-Gordon integration algorithm are chosen as relerr=abserr= 10^{-8} , we get the following errors in the component $\partial \psi / \partial p_1$:

AVM: relerr = $8.06 \cdot 10^{-7}$, abserr = $1.50 \cdot 10^{-8}$, AD+correction: relerr = $2.31 \cdot 10^{-7}$, abserr = $4.30 \cdot 10^{-8}$.

The errors in the other components are similar and omitted here. Both gradients are sufficiently accurate, and both methods can be used, for example, for the sensitivity analysis of multibody systems. Also, for all other components, correct results are computed with both methods. In general, the computation of the adjoint variable gradients is more efficient, but they are based on a hand-coded, highly optimized algorithm whose implementation took man-years, while the AD-generated code is fairly simple to create and requires (including the result verification) much less time, at the expense of a less efficient execution.

We note, however, that, while usually even big codes can be run through ADIFOR within a few days, the time for the verification of the gradients can be much higher. Unless both the algorithm and its implementation are well understood, one should check the results carefully. This need for verification is not due to potential weaknesses of AD tools, but due to the fact that AD differentiates an algorithm without any knowledge of the mathematics that underlied the algorithm. This is both a strength, as programs of arbitrary size can be handled (ADIFOR has successfully differentiated codes of 120,000 lines in length, and produced the desired results), but also a potential weakness, as the discussions so far have shown. Thus, AD is no "silver bullet," but we believe that, at least from our experience in multibody optimization, it does substantially ease the effort required for derivative computation while delivering acceptable performance. Finite difference approximations are in our experience inacceptable due to their lack of efficiency and reliability. The other alternative, the development of specialized code for the gradient computation, is costly in terms of human effort, but can be justified when maximal efficiency is a major design goal, for example, in the development of the adjoint variable method. It is worth emphasizing though, that the implementation of the adjoint variable method is closely tied to a particular integration scheme. It is cumbersome to exchange the integration algorithm, whereas AD techniques allow integration schemes to be substituted quite easily.

5 Conclusions

From the described investigations, we can summarize our conclusions in three groups:

• The numerical behavior of the criteria and the gradient computation must be studied carefully. It is not obvious, for example, that the stepsize control is determined only by the state variables required for the criteria computation and, therefore, the errors introduced in the state variables for the gradients may be bigger than the prescribed error bounds for the state variables for the criteria. Similar issues arise in the context of applying AD to iterative solvers [12]. In our experience, this behavior is acceptable in many practical applications because for optimization purposes, for example, it is often sufficient to compute the gradients with lesser accuracy than the criteria.

- The formulation of the gradient equations for the actual computation performed (versus their counterpart in the world of continuous mathematics), such as (7) may be subtle. For example, the formulation of implicit, parameter-dependent final conditions instead of a fixed final time would introduce new dependencies of the criteria from the intermediate time steps. Also, time discretization and stepsize control are likely not the only influencing factors from numerics. Features such as variable-order polynomial interpolations and projections depend also on the input quantities, become assigned gradients from AD tools and therefore influence the finally computed gradients. Corrections of AD-computed gradients are required to arrive at the mathematically desired results, and the remarks given here for the adaptive time discretization thus may need to be extended to handle other auxiliary variables in an algorithm.
- Despite the fact that the application of plain AD often yields the right results, the inclusion of expert knowledge can highly improve the performance and numerical behavior. If, in our example, the differentiation of the stepsize control in the AD-generated code is switched off, we can compute more efficient and still correct gradients. However, these modifications require a lot of knowledge about the problem and the gradient computation. Thus, even if AD tools provide annotation capabilities that allow a user to treat certain variables as constant with respect to differentiation, one still needs to be careful not to miss any dependencies.

Thus, while the work presented here allowed us to obtain the desired derivatives from an algorithm relevant for practical problems such as the Shampine-Gordon algorithm, we may definitely not conclude that the a posteriori correction of (11) is sufficient for all other integration algorithms as well. AD tools such as ADIFOR or ADIC [5] allow the differentiation of arbitrary complex codes, but for each of them, one must decide whether and which modifications or a posteriori corrections are required to obtain correct results. General rules are hardly possible, but we expect that the work presented here will cover a fair number of cases. Moreover, the work helps to sharpen the users eyes for other possible sources of "errors" arising from the discrepancy between the derivatives of the integration algorithm and the derivatives of the solution that is being approximated by this algorithm.

Acknowledgments

C. B. was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Computational and Technology Research, U.S. Department of Energy, under contract W-31-109-Eng-38. This work was partially completed during P. E is visit to the Department of

This work was partially completed during P. E.'s visit to the Department of Mechanical Engineering at the University of California at Berkeley supported by the German Research Council (DFG) under grant EB195/1-1. Lastly, we thank the anonymous referee for pointing out a major oversight in the initial draft.

References

- M. Berz, C. Bischof, G. Corliss, and A. Griewank (eds.). Computational Differentiation: Techniques, Applications, and Tools. SIAM, Philadelphia, 1996.
- [2] D. Bestle and P. Eberhard. Analyzing and Optimizing Multibody Systems. Mechanics of Structures and Machines, 20(1):67-92, 1992.
- [3] C. Bischof, A. Carle, P. Khademi, and A. Mauer. ADIFOR 2.0: Automatic Differentiation of Fortran77 Programs, *IEEE Computational Science & Engineering*, 3(3):18-32, 1996.
- [4] C. Bischof and M. Haghighat. On Hierarchial Differentiation. In [1], pp. 83-94.
- [5] C. Bischof, L. Roh, and A. Mauer. ADIC An Extensible Automatic Differentiation Tool for ANSI-C. Preprint ANL/MCS-P626-1196, Argonne National Laboratory, 1996 (to appear in Software Practice and Experience).
- [6] J. Butcher. The Numerical Analysis of Ordinary Differential Equations (Runge-Kutta and General Linear Methods). John Wiley and Sons, New York, 1987.
- [7] J. Cash and A. Karp. A Variable Order Runge-Kutta Method for Initial Value Problems with Rapidly Varying Right-Hand Sides. ACM Transactions on Mathematical Software, 16(3):201-222, 1990.
- [8] E. Cramer, P. Frank, G. Shubin, J. Dennis, and L. Michael. Problem Formulations for Multidisciplinary Optimization. SIAM J. Optimization, 4(4):754-776, 1994.
- [9] P. Eberhard. Zur Mehrkriterienoptimierung von Mehrkörpersystemen. VDI-Fortschritt-Berichte, 11, No. 227. VDI-Verlag, Düsseldorf, 1996.
- [10] A. Griewank. On Automatic Differentiation. In Mathematical Programming: Recent Developments and Applications, pp. 83-108. Kluwer Academic Publishers, Amsterdam, 1989.
- [11] A. Griewank and S. Reese. On the Calculation of Jacobian Matrices by the Markowitz Rule. In A. Griewank and G. Corliss (eds.), Automatic Differentiation of Algorithms: Theory, Implementation, and Application, pp. 126-135. SIAM, Philadelphia, 1991.
- [12] A. Griewank, C. Bischof, G. Corliss, A. Carle, and K. Williamson. Derivative Convergence of Iterative Equations Solvers. Optimization Methods and Software, 2:321-355, 1993.
- [13] E. Hairer, S. Norsett, and G. Wanner. Solving Ordinary Differential Equations I: Nonstiff Problems. Springer, New York, 1987.
- [14] E. Hairer and G. Wanner. Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems. Springer, New York, 1991.
- [15] A. Hindmarsh. ODEPACK, a Systematized Collection of ODE Solvers. In R. Stepleman et al. (eds.), *Scientific Computing*, pp. 55-64, IMACS/North-Holland Publishing Company, Amsterdam, 1983.
- [16] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. Numerical Recipes in Fortran: The Art of Scientific Computing. 2nd edition, Cambridge University Press, 1992.

- [17] L. Rall. Automatic Differentiation: Techniques and Applications. Springer, Berlin, 1981.
- [18] A. Sandu, G. Carmichael, and A. Potra. Sensitivity Analysis for Atmospheric Chemistry Models via Automatic Differentiation. Atmospheric Environment, 31(3):475-489, 1997.
- [19] G. Sewell. The Numerical Solution of Ordinary and Partial Differential Equations. Academic Press, San Diego, 1988.
- [20] L. Shampine and M. Gordon. Computer Solution of Ordinary Differential Equations. Freeman, San Francisco, 1975.
- [21] J. Stoer and R. Bulirsch. Introduction to Numerical Analysis. Springer, New York, 1980.