Adjoint Differentiation
of a Structural Dynamics Solver⋆

Mohamed Tadjouddine[1], Shaun A. Forth[1], and Andy J. Keane[2]

[1] Applied Mathematics & Operational Research, ESD, Cranfield University
(RMCS Shrivenham), Swindon, UK
{M.Tadjouddine,S.A.Forth}@cranfield.ac.uk
[2] Computational Engineering and Design Group, School of Engineering Sciences,
University of Southampton, UK
andy.keane@soton.ac.uk

Summary. The design of a satellite boom using passive vibration control by Keane
[J. of Sound and Vibration, 1995, 185(3), 441–453] has previously been carried out
using an energy function of the design geometry aimed at minimising mechanical
noise and vibrations. To minimise this cost function, a Genetic Algorithm (GA) was
used, enabling modification of the initial geometry for a better design. To improve
efficiency, it is proposed to couple the GA with a local search method involving the
gradient of the cost function. In this paper, we detail the generation of an adjoint
solver by automatic differentiation via ADIFOR 3.0. This has resulted in a gradient
code that runs in 7.4 times the time of the function evaluation. This should reduce
the rather time-consuming process (over 10 CPU days by using parallel processing)
of the GA optimiser for this problem.

Key words: Reverse mode AD, hybrid GA-local search, structural dynamics, per-
formance, ADIFOR

1 Introduction

In space missions, lightweight cantilever structures are often used to suspend scien-
tific instruments, such as antenna, a few metres away from the satellite. An example
of this kind of structure is the satellite boom shown in Fig. 1. Vibrations can be
transmitted through the structure from satellite to the instrument. Such vibrations
can damage the boom structure or prevent it from being used for its intended pur-
pose. To ensure correct functioning of the instrument, the vibrations or mechanical
noise through the structure must be kept at tolerable levels. Typically, the structure
is excited by a point transverse force near an end beam, and the energy level is mea-
sured at the opposite end beam. To minimise vibrations and noise, the geometry of
the structure is modified to reduce the frequency average response of the satellite
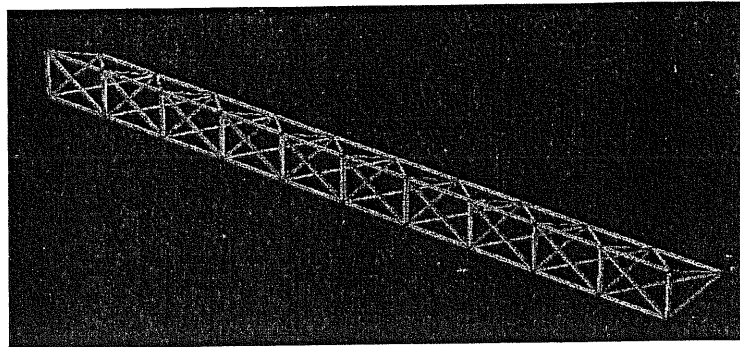boom. This is known as *passive vibration control* [296].

Fig. 1. Initial geometry of the satellite boom.

In this paper, we consider the satellite boom of Fig. 1 described in Sect. 2.1 and previously studied in [298, 384]. The structural dynamics of the three-dimensional satellite boom are modelled by a Fortran computer code named BEAM3D [476, 477] using receptance theory, whereby the behaviour of the global structure is predicted from the Green functions of the individual components, evaluated as summations over their mode shapes.

In previous work [298, 384], the minimisation of the frequency average response (in the range 150–250 Hz) at the end beam was carried out to find a superior design or geometry. For that purpose, a Genetic Algorithm (GA) was used. GAs are known to work for fairly large cost with a good chance of finding the global minimum. Generally, GAs do not require the gradient of the cost function. However, the application of GAs in large scale industrial applications is limited due to the large number of expensive evaluations of the cost function. For our application, the first 10 generations for a population size of 100 took over 10 days to complete using parallel processing [298].

Attention has now shifted to a hybrid genetic algorithm-local search approach combining Darwinian and Lamarckian evolution models. Darwinian evolution, based on "natural selection" considers that the most fit individuals are likely to survive. Lamarckian evolution takes the view that individuals may improve within their environment as they become adapted to it and that the resulting changes are inheritable. Consequently, the genotype of an improved individual is forced to reflect the result of such an improvement by replacing the individual into the population for reproduction [417]. In our application, the local search is carried out using a gradient method, which requires the calculation of the gradient of a cost function $F : \mathbb{R}^{90} \to \mathbb{R}$. Here, the 90 inputs are the coordinates of the joint positions in the design geometry and represent the independent variables. The computation of transmitted power $F(\mathbf{x})$ involves complex variable calculations, which are handled by ADIFOR 3.0 [96] since the dependents and independents are real values [445].

In theory, the gradient ∇F of such a function is cheaply calculated using the reverse mode since the cost of the gradient is independent of the number of the independents and is bounded above by a small factor of the cost of evaluating the function [225]. In practice, large memory requirement may prohibit use of the

adjoint code (reverse AD generated code). This paper details the differentiation of the BEAM3D code by the ADIFOR, the successor of the AD tool ADIFOR 2.0 [57]. ADIFOR employed in reverse mode produces an adjoint code which, after being tuned manually for performance enhancement, calculated the function and its gradient in 7.4 times the CPU time required for its function evaluation. Moreover, the adjoint code runs 12.6 times faster than one-sided finite-differencing (FD) on a Sun Blade 1000 machine with 1200 MHz CPU, 8 MB external cache and 2 GB RAM.

2 Optimisation of the Boom Structure

We aim at minimising vibrations through the structure represented in Fig. 1. There are at least three ways to achieve this: increasing the mass of elements or coating elements with damping material, using active anti-vibration to cancel unwanted vibrations; and as considered here, modifying the geometry of the structure to filter and reflect the vibrations.

2.1 The Initial Geometry

The initial boom structure to be optimised is three-dimensional and composed of 90 Euler-Bernouilli beams each having the same properties per unit length. Because the structure is used to mount a scientific instrument away from a space satellite, the length of the boom structure must be chosen within reasonable limits. Typical values of the aluminium were used for the physical properties of the beams. The bay length is 45 cm, and the overall length of the boom structure is 4.5 m.

The beams were arranged in a regular manner along the XYZ axes so that the YZ cross-section of the boom structure formed an equilateral triangle. The three joints at the left hand end of the structure were fixed, i.e., they were clamped to prevent motion. The beams were connected together with 30 free joints. Geometric constraints were used to avoid beams overlapping or becoming extremely long. The free joints were kept within fixed distances of their original positions. The connectivity of the diagonal beams was chosen so that a maximum of six beams met at any one joint.

Typically, the structure is excited by a point transverse force applied to a left hand end beam of the structure. The vibrational energy level is calculated at a right hand end beam using receptance methods [296]. The optimisation aims at minimising the vibrations by minimising the frequency averaged response in the range $150 - 250$ Hz. For that purpose, the optimiser is allowed to modify the geometry of the satellite boom by changing the coordinates of the 30 free joints in the structure.

2.2 An Optimised Geometry

A GA from the optimisation software package [297] was used to generate an optimised boom geometry by improving the frequency response curve. The principles of a GA can be found in [217]. In short, GAs work on the premise that a population of competing individuals can be combined to produce improved individuals. They mimic "natural" selection, or Darwinian evolution. The number of generations and their population size are usually chosen in advance. Common operations are:

- selection: whereby the fittest individuals are chosen to "inter-breed" and pass their attributes to their offspring.
- crossover: where random portions of two of the most fit individuals are combined to form a new individual.
- mutation: where small changes are introduced to one individual at a time.

As reported in [384], an optimised design geometry was obtained using an objective function F set to be the square root of the sum of velocity squared in the X, Y, Z directions for the end three joints labelled $31, 32$ and 33,

$$F = \int_{\text{freq}} \sum_{j=31}^{33} \sqrt{V_{xj}^2 + V_{yj}^2 + V_{zj}^2} \ .$$

A run of the GA for 10 generations and a population size of 300, gave the novel design geometry of the boom structure shown in Fig. 2. However, GAs applied to large-scale optimisation problems can take CPU days even using parallel processing [298]. To enhance their performance, they may be combined with local search methods.
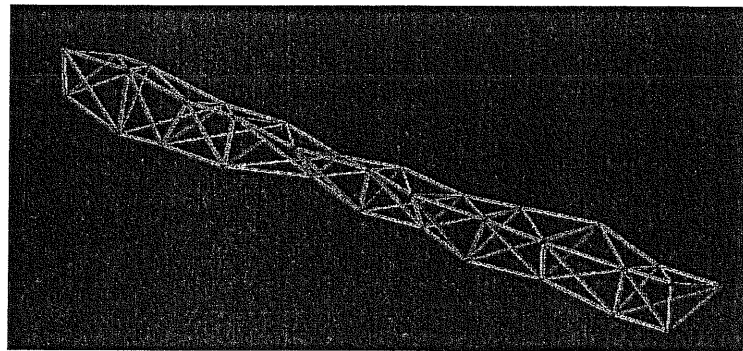

Fig. 2. An optimised geometry of the satellite boom.

2.3 Coupling GAs and Local Search Methods

For optimizing the design of the satellite boom of Fig. 1, a GA coupled with a local search method based on gradient descent should outperform a stand-alone GA. We aim to make them efficient by taking advantage of the accuracy and efficiency of gradient calculation by reverse mode AD.

In essence, the hybridised GA-Local Search method performs similar steps to that of the GA except that each individual of the population is locally improved using a local search method, here steepest descents, following the meta-Lamarkian learning approach as shown in Fig. 3 and detailed in [417]. In Fig. 3, the while loop is executed until either t exceeds some maximum number of generations t_{max} or convergence is detected. We now detail the differentiation of the BEAM3D code to enable the gradient descent method of the hybrid algorithm.

```
t = 0
Initialise a GA population P(t) = {x1, x2, ..., x^m}
While ( EndCondition is not satisfied )
      Evaluate Fitness(P(t)) giving F(x1), F(x2), ..., F(x^m)
      For each individual x^i ∈ P(t)
            Improve x^i using the Gradient Descent method GD(x^i)
            Replace x^i by the improved x_new^i = GD(x^i) in P(t)
      EndFor
      Generate P(t + 1) from the x_new^i by using standard
      GA operations (Selection, Mutation, or Crossover).
      t = t + 1
EndDo
```

Fig. 3. Hybrid GA-Gradient Descent Method.

3 Differentiation of the BEAM3D Code

The BEAM3D code, as sketched in Fig. 4, starts by reading in data from files representing the boom geometry and certain properties of each beam. Given extra information such as the range of frequencies over which to solve, the number of data points within the specified frequency range, the joint numbers at which to calculate the energies, and the number of axial[torsional] and transverse modes in the modal summations for the Green functions; the program builds up a linear complex system $\mathbf{A}\mathbf{f} = \mathbf{b}$ for nodal forces $\mathbf{f}$ and solves it for each frequency. An averaged energy function is calculated at the specified end beam.

```
Read in:
      No. of beams, beam properties, and list of connections
      Frequency range [ω_min, ω_max, N]  (N ≡ No. of frequencies)
      Coordinates of beam ends x_{n,j}
Calculate some geometric information
Initialise F = 0 (integral of power)
Δω = (ω_max − ω_min)/(N − 1)
For k = 1, N
      ω = ω_min + (k − 1) * Δω
      Assemble:
            Green Function Matrix A(x, iω)
            r.h.s. forcing b(x, iω)
      Solve A(x, iω)f = b(x, iω) (LAPACK) - 50% of CPU time
      Obtain displacement D_{n,j} at ends of beam n
      Obtain power P = ½ Re(f_j · D_{n,j})
      Update integral F = F + P2 * Δω
End For
```

Fig. 4. Schematic of the BEAM3D code.

We aim to calculate the sensitivities of the energy function with respect to the coordinates of the free joints. This represents a gradient calculation with 90 independent variables. Prior to differentiation, the code was restructured so all reading of data is done outside the subroutines to be differentiated. Furthermore, to allow the code to be processed by ADIFOR, the code was rewritten according to the Fortran 77 standard. Actually, the original code contained (non-standard) language extensions in the form of structures defined as follows:

```
        STRUCTURE /PROPERTY/
                INTEGER ID
                CHARACTER*6 ENDCON
                DOUBLE PRECISION ANGLE(3,3)
                DOUBLE PRECISION LENG
                ...
                COMPLEX *16  FM2
        END STRUCTURE
        RECORD /PROPERTY/ BEAM(150)
```

This structure is replaced using arrays corresponding to the components of the structure. The restructured code contains the following array declarations:

```
        INTEGER BEAM_ID(150)
        CHARACTER*6 BEAM_ENDCON(150)
        DOUBLE PRECISION BEAM_ANGLE(3,3,150)
        DOUBLE PRECISION BEAM_LENG(150)
        ...
        COMPLEX*16 BEAM_FM2(150)
```

A sed [153] script was written to replace any instance BEAM(I).X(K,J), where X represents any component of the structure, by the array element BEAM_X(K,J,I). If X is a scalar variable, obviously no indices are used. The resulting computer code is differentiated using FD, and AD via ADIFOR.

3.1 Initial Differentiation

We first computed a single directional derivative $\dot{y} = \nabla F(\mathbf{x})\dot{\mathbf{x}}$ for a random direction $\dot{\mathbf{x}}$, by using one sided FD, AD in forward mode, and a single adjoint $\bar{\mathbf{x}} = \nabla F(\mathbf{x})^T \bar{\mathbf{y}}$ for $\bar{\mathbf{y}} = 1$ via reverse mode AD. By definition of the adjoint operator, we have $\bar{\mathbf{y}}\dot{\mathbf{y}} \equiv \bar{\mathbf{x}}\dot{\mathbf{x}}$, which allows us to validate the results of the differentiation. The initial ADIFOR generated codes gave incorrect results inconsistent with those from FD, caused partly by non-differentiable statements in the code for the function F.

3.2 Dealing With Non-Differentiability

A major assumption in AD is that the function F to be differentiated is composed of elemental functions ϕ that are continuously differentiable on their open domains [225]. At a point on the boundary of an open domain, F is continuous, but ∇F may jump to a finite value or even infinity. This is important when the computer code that represents the function contains branches, some kink functions

(e.g., `abs`), or inverse functions (e.g., `sqrt`, `arctan`, or `arccos`). To compute reliable derivatives, such pathological cases must be handled correctly. It is known that these cases can be tackled by calculating derivatives in a given direction [225]. Insights or knowledge of the computer code can also be exploited. The BEAM3D code contains at least two types of non-differentiability.

```
l2 = datan2(xdiff,zdiff)
m2 = dsqrt(xdiff*xdiff+zdiff*zdiff)/beam_leng(i)
sgn1 = -1.0d0
sgn2 = -1.0d0
sgn3 = -1.0d0
if (xdiff.lt.0.0d0) sgn1 = -sgn1
if (ydiff.gt.0.0d0) sgn2 = -sgn2
if (zdiff.lt.0.0d0) sgn3 = -sgn3
yor(1,i) = sgn1*dabs(dsin(dacos(m2))*dsin(l2))
yor(2,i) = sgn2*dabs(m2)
yor(3,i) = sgn3*dabs(dsin(dacos(m2))*dcos(l2))
```

Fig. 5. A code fragment that is non-differentiable.

The first type of non-differentiability is due to the presence of the functions `arccos` and `abs`. ADIFOR allows us to locate possible non-differentiable points by generating the derivative code with the *Exception Handling* enabled [96]. On running this code, warnings were raised concerning the functions `arccos` and `abs`. The part of the code, containing such anomalies is shown in Fig. 5. We then used algebra and trigonometric formula to rewrite some of the algebraic expressions containing (`arccos`, `abs`, `sin` and `tan`) as in Fig. 6.

This transformation resulted in equivalent expressions calculating the same values but differentiable in the vicinity of their arguments.

```
myt1 = ydiff/beam_leng(i)
myt2 = dsqrt(xdiff*xdiff+zdiff*zdiff)
sgn1 = -1.0d0
sgn2 = -1.0d0
sgn3 = -1.0d0
if (xdiff.lt.0.0d0) sgn1 = -sgn1
if (ydiff.gt.0.0d0) sgn2 = -sgn2
if (zdiff.lt.0.0d0) sgn3 = -sgn3
yor(1,i) = sgn1*myt1*xdiff/myt2
yor(2,i) = sgn2*myt2/beam_leng(i)
yor(3,i) = sgn3*myt1*zdiff/myt2
```

Fig. 6. An equivalent but differentiable version of the code fragment of Fig. 5.

The second type of non-differentiability encountered in BEAM3D was due to a branching construct illustrated by the code fragment of Fig. 7. As xdiff and ydiff are active variables, the differentiation of this code fragment gave point-valued derivatives that prevented the function **F** from being differentiable. Such branches represent constraints on the design geometry and, in our case, may be safely removed.

```
if (xdiff.eq.0.0 .and. ydiff.eq.0.0)
then
    yor(1,i) = zdiff/beam_leng(i)
    yor(2,i) = 0.0
    yor(3,i) = 0.0
else ....
```

Fig. 7. A code fragment testing whether an active variable is zero.

Finally, the complex linear solver **Af**=**b** employs the LAPACK routine zgesv [7]. Differentiating the LAPACK source code routines for zgesv using an AD tool without taking account insights into the nature of the linear solver would give inefficient code. Mechanical generation of the zgesv derivative by ADIFOR gave not only inefficient code but also results inconsistent with FD. Therefore, we hand-coded its derivative as described in Sect. 3.3.

3.3 Complex Linear Solver

Instead of using ADIFOR to differentiate the complex linear solve,

$$\mathbf{Af} = \mathbf{b} , \tag{1}$$

of the LAPACK routine zgesv, we instead use hand-coding for both forward and reverse mode. Differentiating $\mathbf{Af} = \mathbf{b}$, using the matrix-equivalent of the product rule for a single directional derivative, we obtain $\mathbf{A\dot{f}} + \mathbf{\dot{A}f} = \mathbf{\dot{b}}$, and giving the derivatives $\mathbf{\dot{f}}$ by the solution of

$$\mathbf{A\dot{f}} = \mathbf{\dot{b}} - \mathbf{\dot{A}f} . \tag{2}$$

In the forward mode, we may re-use the LU-decomposition of **A** to solve efficiently for the derivatives $\mathbf{\dot{f}}$. The following procedure is used:

1. Perform an **LU** decomposition of the matrix **A**
2. Solve **Af** = **b**
3. Form $\mathbf{b}_{new} = \mathbf{\dot{b}} - \mathbf{\dot{A}f}$
4. Re-use LU-decomposition to solve $\mathbf{A\dot{f}} = \mathbf{b}_{new}$

$\mathbf{\dot{A}}$ and $\mathbf{\dot{b}}$ are calculated by applying ADIFOR to the Assemble procedure in Fig. 4.

For the reverse mode, deriving the adjoint update corresponding to (1) is more problematic. Defining $\mathbf{C} = \mathbf{A}^{-1}$, we may write

$$\mathbf{\dot{f}} = \mathbf{C\dot{b}} - \mathbf{C\dot{A}f} .$$

Then $\dot{f}_i$, the i^{th} element of $\mathbf{\dot{f}}$, is given by

$$\dot{f}_i = \sum_j c_{ij}\dot{b}_j - \sum_j c_{ij} \sum_k \dot{a}_{jk}f_k , \tag{3}$$

where c_{ij}, $\dot{a}_{jk}$, $\dot{b}_j$, and f_k are the elements of $\mathbf{C}$, $\dot{\mathbf{A}}$, $\dot{\mathbf{b}}$, and $\mathbf{f}$, respectively. Now we use the identity $\bar{\mathbf{y}}\dot{\mathbf{y}} \equiv \bar{\mathbf{x}}\dot{\mathbf{x}}$ for the system $\mathbf{y} = \mathbf{F}(\mathbf{x})$ [225, Equation (3.7)]. In the context of the vector and matrix arguments of the system (1), this identity gives

$$\sum_i \bar{f}_i \dot{f}_i = \sum_i \bar{b}_i \dot{b}_i + \sum_i \sum_j \bar{a}_{ij} \dot{a}_{ij} . \tag{4}$$

From (3), we obtain

$$\sum_i \bar{f}_i \dot{f}_i = \sum_i \bar{f}_i \sum_j c_{ij} \dot{b}_j - \sum_i \bar{f}_i \sum_j c_{ij} \sum_k \dot{a}_{jk} f_k ,$$

by reordering summations,

$$\sum_i \bar{f}_i \dot{f}_i = \sum_j \dot{b}_j \sum_i c_{ij} \bar{f}_i - \sum_j \sum_k \dot{a}_{jk} f_k \sum_i c_{ij} \bar{f}_i ,$$

and by swapping indices (i, j, k) to (k, i, j),

$$\sum_i \bar{f}_i \dot{f}_i = \sum_i \dot{b}_i \sum_j c_{ji} \bar{f}_j - \sum_i \sum_j \dot{a}_{ij} f_j \sum_k c_{ki} \bar{f}_k .$$

Comparing with (4), we see that

$$\bar{b}_i = \sum_j c_{ji} \bar{f}_j ,$$

giving $\bar{\mathbf{b}} = \mathbf{C}^T \bar{\mathbf{f}} = \mathbf{A}^{-T} \bar{\mathbf{f}}$, or that $\bar{\mathbf{b}}$ is the solution of

$$\mathbf{A}^T \bar{\mathbf{b}} = \bar{\mathbf{f}} . \tag{5}$$

Similarly,

$$\bar{a}_{ij} = -f_j \sum_k c_{ki} \bar{f}_k = -f_j \bar{b}_i , \text{ or}$$

$$\bar{\mathbf{A}} = -\bar{\mathbf{b}}\mathbf{f}^T . \tag{6}$$

The adjoint $\bar{\mathbf{b}}$ is updated by solving the linear system (5), while $\bar{\mathbf{A}}$ is updated by adding the right hand side of the equation (6). Adjoint formulae (5) and (6) are equivalent to those given in [520].

Using (5) and (6), we obtain the following procedure for the adjoint of the linear solve:

1. In the forward sweep,
    a) Perform an **LU** decomposition of the matrix **A**,
    b) Store **L** and **U** and the pivot sequence IPIV,
    c) Solve **Af = b**.
2. In the reverse sweep,
    a) Load **L** and **U** and the pivot sequence IPIV,
    b) Solve $\mathbf{A}^T \bar{\mathbf{b}} = \bar{\mathbf{f}}$ for $\bar{\mathbf{b}}$,
    c) Update $\bar{\mathbf{A}} = \bar{\mathbf{A}} - \bar{\mathbf{b}}\mathbf{f}^T$.

Since both **A** and **b** depend on the beam endpoint location **x** (see Fig. 4), the adjoined linear solver must modify their adjoints $\bar{\mathbf{A}}$ and $\bar{\mathbf{b}}$. The adjoint for $\bar{\mathbf{A}}$ is the usual increment, while for $\bar{\mathbf{b}}$ it is an assignment because the LAPACK routine zgesv overwrites b with f. Here, the memory storage is dramatically reduced compared with black-box application of ADIFOR. If **A** is an $N \times N$ matrix, we store only $N2$ complex coefficients instead of $O(N3)$ when ADIFOR tapes all variables on the left of assignment statements in the LU decomposition.

3.4 Initial Results and Validation

After implementing the procedures described in Sect. 3.2 and 3.3 on the ADIFOR generated code, we obtained tangent and adjoint derivative codes that calculate directional derivatives consistent with one-sided FD. The obtained codes were compiled with maximum compile optimisations and run on a Sun Blade 1000 machine. Table 1 shows the results and timings of forward mode AD, reverse mode AD, and one-sided FD for that calculation. These results showed that forward and reverse AD gave the same directional derivative value within roundoff, while the maximum difference with the FD result is around 10^{-6}. This difference is of the order of the square root of the machine relative precision. This validates the AD results as being in agreement with the one-sided FD result.

Table 1. Results for a single directional derivative, timings are in CPU seconds.

Method	$\bar{x}\dot{x}$	$\bar{y}\dot{y}$	CPU$(F, \nabla F)$
FD (1-sided)		0.124578003587	48.7
ADIFOR(fwd)		0.124571139127	54.0
ADIFOR(rev)	0.124571139130		311.5

From Table 1, we see that while the AD reverse mode calculates the gradient in around 5 minutes, one-sided FD and forward AD requires 91 function evaluations and 90 directional derivatives respectively and consequently run times of over 35 minutes. We see that using reverse mode AD can speed up the gradient calculation by a factor of around 7 over FD while giving accurate derivatives. However, the core of the calculation (building the linear system, solving it, and calculating the local energy contribution for each frequency) of the BEAM3D code is an independent loop and therefore can be differentiated in parallel as we now describe.

4 Performance Issues

Usually, after checking that the AD forward and reverse modes agreed with the finite differences, we seek to improve efficiency of the automatically generated code. As shown by the results of Table 1, the reverse mode is superior to the finite differences and forward mode, but it requires a very large amount of memory to run because the tape required 12 GB. By hand coding the adjoint of the linear solver, we reduced the size of the tape to around 6 GB.

Furthermore, the core of the calculation of the BEAM3D code is carried out in a parallel loop, which is the loop over k in Fig. 4. Because the iterations of such a loop are independent, we can run the loop body taping all the required information in just one iteration, then immediately adjoin the body of the loop [254]. This reduced the tape size of the adjoint code to around 0.3 GB. This represents a memory reduction by a factor of 20, the number of extra iterations performed by the parallel loop.

The second row of Table 2 shows that after this optimisation, the ratio between the gradient calculation and the function is 7.4. It also shows a speed up factor of 12.6 over the popular one-sided FD method.

Table 2. CPU Timings (in Seconds) on a SUN Blade 1000, UltraSparcIII.

Method	CPU$(F, \nabla F)$	CPU$(F, \nabla F)$/CPU(F)
ADIFOR(rev.)	311.5	13.3
ADIFOR(rev.,par.)	174.7	7.4
FD (1-sided)	2215.9	93.1

5 Conclusions

ADIFOR allowed us to build an adjoint for a code that makes extensive use of complex variable arithmetic to accurately calculate the gradient of a cost function. The adjoint code requires only 7.4 times the CPU time of the original function code, and the memory requirement for taping is a modest 0.3 GB. It also runs 12.6 times faster than calculating the gradient using one-sided finite differencing.

Future work is planned to compare the performance of gradient calculation using both ADIFOR [96] and TAF [163] capabilities. The design optimisation of the lightweight cantilever structure will be carried out using the meta-Lamarckian learning strategy [417], which efficiently combines GAs with local search methods. The reduction in computational time of the gradient calculation will be of great benefit in allowing the meta-Lamarckian algorithm to be used to optimise the design of boom structures.

Acknowledgements

Lecture Notes in Computational
Science and Engineering

50

$$\frac{\mathrm{d}\phi_t}{\mathrm{d}v_s} = \sum_{s \overset{p}{\rightsquigarrow} t} \prod_{(i,j)\in p} \frac{\partial\phi_j}{\partial v_i}$$
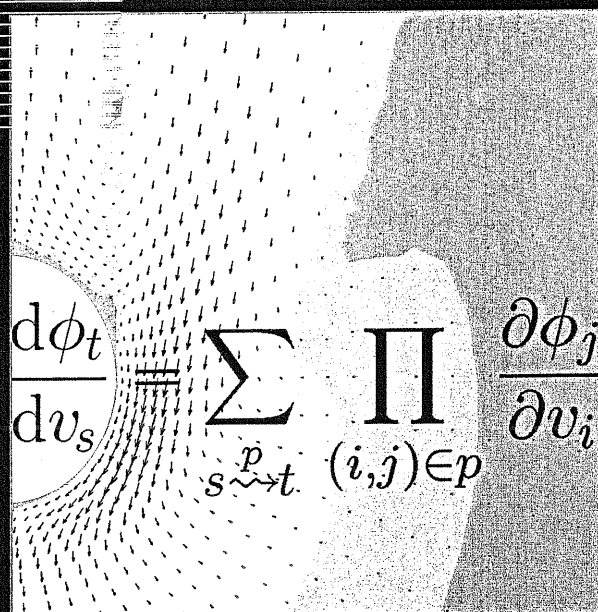
Martin Bücker
George Corliss
Paul Hovland
Uwe Naumann
Boyana Norris
Editors

Automatic Differentiation: Applications, Theory, and Implementations

Springer

Editors

Martin Bücker

Institute for Scientific Computing
RWTH Aachen University
D-52056 Aachen, Germany
email: buecker@sc.rwth-aachen.de

Uwe Naumann

Software and Tools
for Computational Engineering
RWTH Aachen University
D-52056 Aachen, Germany
email: naumann@stce.rwth-aachen.de

George Corliss

Department of Electrical
and Computer Engineering
Marquette University
1515 W. Wisconsin Avenue
P.O. Box 1881
Milwaukee, WI 53201-1881, USA
email: george.corliss@marquette.edu

Paul Hovland
Boyana Norris

Mathematics and Computer Science Division
Argonne National Laboratory
9700 S Cass Ave
Argonne, IL 60439, USA
email: hovland@mcs.anl.gov
      norris@mcs.anl.gov