ADOL-C: [1]

A Package for the Automatic Differentiation of Algorithms Written in C/C++

Version 1.10.0, July 2005

Andrea Walther[2], Andreas Kowarz[2], and Andreas Griewank[3]

Abstract

The C++ package ADOL-C described here facilitates the evaluation of first and higher derivatives of vector functions that are defined by computer programs written in C or C++. The resulting derivative evaluation routines may be called from C, C++, Fortran, or any other language that can be linked with C.

The numerical values of derivative vectors are obtained free of truncation errors at a small multiple of the run time and random access memory required by the given function evaluation program. Derivative matrices are obtained by columns, by rows or in sparse format. For solution curves defined by ordinary differential equations, special routines are provided that evaluate the Taylor coefficient vectors and their Jacobians with respect to the current state vector. For explicitly or implicitly defined functions derivative tensors are obtained with a complexity that grows only quadratically in their degree. The derivative calculations involve a possibly substantial but always predictable amount of data. Since the data is accessed strictly sequentially it can be automatically paged out to external files.

Keywords: Computational Differentiation, Automatic Differentiation, Chain Rule, Overloading, Taylor Coefficients, Gradients, Hessians, Forward Mode, Reverse Mode, Implicit Function Differentiation, Inverse Function Differentiation

Abbreviated title: Automatic differentiation by overloading in C++

[1]The development of earlier versions was supported by the Office of Scientific Computing, U.S. Department of Energy, the NSF, and the Deutsche Forschungsgemeinschaft. During the development of the current version the co-authors Walther and Kowarz are supported by the grant Wa 1607/2-1 of the Deutsche Forschungsgemeinschaft

[2]Institute of Scientific Computing, Technical University Dresden, D-01062, Germany

[3]Department of Mathematics, Humboldt-Universität zu Berlin, D-10099, Germany

Contents

1 Introduction: Differentiation of Algorithms **4**

2 Preparing a Section of C or C++ Code for Differentiation **6**

2.1 Declaring Active Variables 6
2.2 Marking Active Sections 6
2.3 Selecting Independent and Dependent Variables 7
2.4 A Subprogram as an Active Section 8
2.5 Overloaded Operators and Functions 9
2.6 Reusing the Tape for Arbitrary Input Values 11
2.7 Conditional Assignments 12
2.8 Step-by-Step Modification Procedure 13

3 Numbering the Tapes and Controlling the Buffer **14**

3.1 Examining the Tape and Predicting Storage Requirements 15
3.2 Customizing ADOL-C 16
3.3 Warnings and Suggestions for Improved Efficiency 17

4 Easy-To-Use Drivers **19**

4.1 Drivers for Optimization and Nonlinear Equations 19
4.2 Drivers for Ordinary Differential Equations 22
4.3 Drivers for Sparse Jacobians and Sparse Hessians 24
4.4 Higher Derivative Tensors 29
4.5 Derivatives of Implicit and Inverse Functions 31

5 Basic Drivers for the Forward and Reverse Mode **33**

6 Overloaded Forward and Reverse Calls **39**

6.1 The Scalar Case 39
6.2 The Vector Case 40
6.3 Dependence Analysis 42

7 Tapeless forward differentiation in ADOL-C **44**

7.1 Modifying the Source Code 44

7.2 Compiling and Linking the Source Code 49

7.3 Concluding Remarks for the Tapeless Forward Mode Variant 49

8 Installing and Using ADOL-C **51**

8.1 Generating the ADOL-C Library 51

8.2 Compiling and Linking the Example Programs 52

8.3 Description of Important Header Files 52

8.4 Compiling and Linking C/C++ Programs 53

8.5 Adding Quadratures as Special Functions 53

9 Example Codes **55**

9.1 Speelpenning's Example (`speelpenning.cpp`) 55

9.2 Power Example (`powexam.cpp`) 57

9.3 Determinant Example (`detexam.cpp`) 58

9.4 Ordinary Differential Equation Example (`odexam.cpp`) 60

1 Introduction: Differentiation of Algorithms

Most nonlinear computations require the evaluation of first and higher derivatives of vector functions with m components in n real or complex variables. This requirement arises particularly in optimization, nonlinear equation solving, numerical studies of bifurcation, and the solution of nonlinear differential or integral equations. Often these functions are defined by sequential evaluation procedures involving many intermediate variables. By eliminating the intermediate variables symbolically, it is theoretically always possible to express the m dependent variables directly in terms of the n independent variables. Typically, however, the attempt results in unwieldy algebraic formulae, if it can be completed at all. Symbolic differentiation of the resulting formulae will usually exacerbate this problem of *expression swell* and often entails the repeated evaluation of common expressions.

An obvious way to avoid such redundant calculations is to apply an optimizing compiler to the source code that can be generated from the symbolic representation of the derivatives in question. Exactly this approach was investigated by Bert Speelpenning, a student of Bill Gear, during his Ph.D. research [23] at the University of Illinois from 1977 to 1980. Eventually he realized that at least in the cases $n = 1$ and $m = 1$, the most efficient code for the evaluation of derivatives can be obtained directly from that for the evaluation of the underlying vector function. In other words, he advocated the differentiation of evaluation algorithms rather than formulae. In his thesis he made the particularly striking observation that the gradient of a scalar-valued function (i.e. $m = 1$) can always be obtained for no more than five times the operation count of evaluating the function itself. This bound is completely independent of n, the number of independent variables, and allows the row-wise computation of Jacobians for at most $5\,m$ times the effort of evaluating the underlying vector function.

If m, the number of component functions, is larger than n, Jacobians can be obtained more cheaply column by column through propagating gradients forward. This classical technique of automatic differentiation goes back at least to Wengert [26] and was later popularized by Rall [21]. It was noted in [10] that in general neither the row-by-row nor the column-by-column method is optimal for the calculation of Jacobians. The potentially more efficient alternatives, however, require some combinatorial optimization and involve large data structures that are not necessarily accessed sequentially [13]. Therefore, the package ADOL-C described here was written primarily for the evaluation of derivative vectors (e.g. rows or columns of Jacobians). This approach also simplifies parameter passing between subroutines and calls from different computer languages.

The *reverse propagation of gradients* employed by Speelpenning is closely related to the *adjoint sensitivity analysis* for differential equations, which has been used at least since the late sixties, especially in nuclear engineering [5],[6], weather forecasting [24], and neural networks [27]. The discrete analog used here was apparently first discovered in the early seventies by Ostrovskii et al. [20] and Linnainmaa [19] in the context of rounding error estimates. Since then, there have been numerous re-discoveries and various software imple-

mentations. Currently, there exist at least four precompilers for automatic differentiation, namely ADIFOR[4], OpenAD[5], TAPENADE[6], and TAF[7].

In contrast to precompilation, overloading requires only minor modifications of the user's evaluation program and does not generate intermediate source code. Our package ADOL-C utilizes overloading in C++, but the user has to know only C. The acronym stands for **A**utomatic **D**ifferentiation by **O**ver**L**oading in **C**++. As starting points to retrieve further information on techniques and application of automatic differentiation, as well as on other AD tools, we refer to the conference proceedings [12], [2], [8], and the web page `http://www.autodiff.org`.

ADOL-C facilitates the simultaneous evaluation of arbitrarily high directional derivatives and the gradients of these Taylor coefficients with respect to all independent variables. Relative to the cost of evaluating the underlying function, the cost for evaluating any such scalar-vector pair grows as the square of the degree of the derivative but is still completely independent of the numbers m and n.

For the reverse propagation of derivatives, the whole execution trace of the original evaluation program must be recorded, unless it is recalculated in pieces as advocated in [11, 15]. In ADOL-C, this potentially very large data set is written first into a buffer array and later into a file if the buffer is full or if the user wishes a permanent record of the execution trace. In either case, we will refer to the recorded data as the *tape*. The user may create several tapes in several named arrays or files. During subsequent derivative evaluations, tapes are always accessed strictly sequentially, so that they can be paged in and out to disk without significant runtime penalties. If written into a file, the tapes are self-contained and can be used by other Fortran, C or C++ programs.

This manual is organized as follows. Section 2 explains the modifications required to convert undifferentiated code to code that compiles with ADOL-C. Section 3 covers aspects of the tape of recorded data that ADOL-C uses to evaluate arbitrarily high order derivatives. The discussion includes storage requirements and the tailoring of certain tape characteristics to fit specific user needs. Descriptions of easy-to-use drivers for a convenient derivative evaluation are contained in Section 4. Section 5 offers a more mathematical characterization of the different modes of AD to compute derivatives. At the same time, the corresponding drivers of ADOL-C are explained. The overloaded derivative evaluation routines using the forward and the reverse mode of AD are explained in Section 6. The tapeless forward mode is presented in Section 7. Section 8 details the installation and use of the ADOL-C package. Finally, Section 9 furnishes some example programs that incorporate the ADOL-C package to evaluate first and higher-order derivatives. These and other examples are distributed with the ADOL-C source code. The user should simply refer to them if the more abstract and general descriptions of ADOL-C provided in this document do not suffice.

[4]Contact: Paul Hovland, ANL-MCS, USA, e-mail: hovland@mcs.anl.gov
[5]Contact: Uwe Naumann, RWTH Aachen, Germany, e-mail: naumann@stce.rwth-aachen.de
[6]Contact: Laurent Hascoët, INRIA Sophia-Antipolis, France, e-mail: Laurent.Hascoet@inria.fr
[7]Contact: Ralf Giering, FastOpt GbR, Hamburg, Germany, e-mail: Ralf.Giering@fastopt.de

2 Preparing a Section of C or C++ Code for Differentiation

ADOL-C was designed so that only minimal changes to the users' undifferentiated source code are required. The main modifications concern variable declarations and input/output operations.

2.1 Declaring Active Variables

The key ingredient of automatic differentiation by overloading is the concept of an *active variable*. All variables that may be considered as differentiable quantities at some time during the program execution must be of an active type. ADOL-C uses one active scalar type, called adouble, whose real part is of the standard type double. Typically, one will declare the independent variables and all quantities that directly or indirectly depend on them as *active*. Other variables that do not depend on the independent variables but enter, for example, as parameters, may remain one of the *passive* types double, float, or int. There is no implicit type conversion from adouble to any of these passive types; thus, **failure to declare variables as active when they depend on other active variables will result in a compile-time error message**. In data flow terminology, the set of active variable names must contain all its successors in the dependency graph. All components of indexed arrays must have the same activity status.

The real component of an adouble x can be extracted as x.value(). In particular, such explicit conversions are needed for the standard output procedure printf. The output stream operator $\ll$ is overloaded such that first the real part of an adouble and then the string "(a)" is added to the stream. The input stream operator $\gg$ can be used to assign a constant value to an adouble. Naturally, adoubles may be components of vectors, matrices, and other arrays, as well as members of structures or classes.

The C++ class adouble, its member functions, and the overloaded versions of all arithmetic operations, comparison operators, and most ANSI C functions are contained in the file `<adolc/adouble.cpp>` and its header `<adolc/adouble.h>`. The latter must be included for compilation of all program files containing adoubles and corresponding operations.

2.2 Marking Active Sections

All calculations involving active variables that occur between the void function calls

trace_on(tag,keep) and trace_off(file)

are recorded on a sequential data set called *tape*. Pairs of these function calls can appear anywhere in a C++ program, but they must not overlap. The nonnegative integer argument tag identifies the particular tape for subsequent function or derivative evaluations. Unless

several tapes need to be kept, tag = 0 may be used throughout. The optional integer arguments keep and file will be discussed in Section 3. We will refer to the sequence of statements executed between a particular call to trace_on and the following call to trace_off as an *active section* of the code. The same active section may be entered repeatedly, and one can successively generate several traces on distinct tapes by changing the value of tag. Both functions trace_on and trace_off are prototyped in the header file <adolc/taputil.h>, which is included by the header <adolc/adouble.h> automatically.

Active sections may contain nested or even recursive calls to functions provided by the user. Naturally, their formal and actual parameters must have matching types. In particular, the functions must be compiled with their active variables declared as adoubles and with the header file <adolc/adouble.h> included. Variables of type adouble may be declared outside an active section and need not go out of scope before the end of an active section. It is not necessary – though desirable – that free-store adoubles allocated within an active section be deleted before its completion. The values of all adoubles that exist at the beginning and end of an active section are automatically recorded by trace_on and trace_off, respectively.

2.3 Selecting Independent and Dependent Variables

One or more active variables that are read in or initialized to the values of constants or passive variables must be distinguished as independent variables. Other active variables that are similarly initialized may be considered as temporaries (e.g., a variable that accumulates the partial sums of a scalar product after being initialized to zero). In order to distinguish an active variable x as independent, ADOL-C requires an assignment of the form

$$\mathsf{x} \ll= \mathsf{px} \quad // \ \mathsf{px} \text{ of any passive numeric type} \quad .$$

This special initialization ensures that x.value() = px, and it should precede any other assignment to x. However, x may be reassigned other values subsequently. Similarly, one or more active variables y must be distinguished as dependent by an assignment of the form

$$\mathsf{y} \gg= \mathsf{py} \quad // \ \mathsf{py} \text{ of any passive type} \quad ,$$

which ensures that py = y.value() and should not be succeeded by any other assignment to y. However, a dependent variable y may have been assigned other real values previously, and it could even be an independent variable as well. The derivative values calculated after the completion of an active section always represent **derivatives of the final values of the dependent variables with respect to the initial values of the independent variables**.

The order in which the independent and dependent variables are marked by the $\ll=$ and $\gg=$ statements matters crucially for the subsequent derivative evaluations. However,

these variables do not have to be combined into contiguous vectors. ADOL-C counts the number of independent and dependent variable specifications within each active section and records them in the header of the tape.

2.4 A Subprogram as an Active Section

As a generic example let us consider a C(++) function of the form shown in Figure 1.

```
void eval(int n, int m,          // number of independents and dependents
          double *x,             // independent variable vector
          double *y,             // dependent variable vector
          int *k,                // integer parameters
          double *z)             // real parameters
{                                // beginning of function body
    double t = 0;                // local variable declaration
    for (int i=0; i < n; i++)    // begin of computation
       t += z[i]*x[i];           // continue
    ............                 // continue
    y[m-1] = t/m;                // end of computation
}                                // end of function
```

Figure 1: Generic example of a subprogram to be activated

If eval is to be called from within an active C(++) section with x and y as vectors of adoubles and the other parameters passive, then one merely has to change the type declarations of all variables that depend on x from double or float to adouble. Subsequently, the subprogram must be compiled with the header file `<adolc/adouble.h>` included as described in Section 2.1. Now let us consider the situation when eval is still to be called with integer and real arguments, possibly from a program written in Fortran77, which does not allow overloading.

To automatically compute derivatives of the dependent variables y with respect to the independent variables x, we can make the body of the function into an active section. For example, we may modify the previous program segment as in Figure 2. The renaming and doubling up of the original independent and dependent variable vectors by active counterparts may seem at first a bit clumsy. However, this transformation has the advantage that the calling sequence and the computational part, i.e., where the function is really evaluated, of eval remain completely unaltered. If the temporary variable t had remained a double, the code would not compile, because of a type conflict in the assignment following the declaration. More detailed example codes are listed in Section 9.

```
void eval( int n,m,                  // number of independents and dependents
           double *px,               // independent passive variable vector
           double *py,               // dependent passive variable vector
           int *k,                   // integer parameters
           double *z)                // parameter vector
{                                    // beginning of function body
  short int tag = 0;                 // tape array and/or tape file specifier
  trace_on(tag);                     // start tracing
  adouble *x, *y;                    // declare active variable pointers
  x = new adouble[n];                // declare active independent variables
  y = new adouble[m];                // declare active dependent variables
  for (int i=0; i < n; i++)
      x[i] <<= px[i];                // select independent variables
  adouble t = 0;                     // local variable declaration
  for (int i=0; i < n; i++)          // begin crunch
     t += z[i]*x[i];                 // continue crunch
  .............                      // continue crunch
  .............                      // continue crunch
  y[m-1] = t/m;                      // end crunch as before
  for (int j=0; j < m; j++)
     y[j] >>= py[j];                 // select dependent variables
  delete[] y;                        // delete dependent active variables
  delete[] x;                        // delete independent active variables
  trace_off();                       // complete tape
}                                    // end of function
```

Figure 2: Activated version of the code listed in Figure 1

2.5 Overloaded Operators and Functions

As in the subprogram discussed above, the actual computational statements of a C(++) code need not be altered for the purposes of automatic differentiation. All arithmetic operations, as well as the comparison and assignment operators, are overloaded, so any or all of their operands can be an active variable. An adouble x occurring in a comparison operator is effectively replaced by its real value x.value(). Most functions contained in the ANSI C standard for the math library are overloaded for active arguments. The only exceptions are the non-differentiable functions fmod and modf. Otherwise, legitimate C code in active sections can remain completely unchanged, provided the direct output of active variables is avoided. The rest of this subsection may be skipped by first time users who are not worried about marginal issues of differentiability and efficiency.

The modulus fabs(x) is everywhere Lipschitz continuous but not properly differentiable at the origin, which raises the question of how this exception ought to be handled. Fortunately, one can easily see that fabs(x) and all its compositions with smooth functions are still directionally differentiable. These directional derivatives of arbitrary order can be propagated in the forward mode without any ambiguity. In other words, the forward mode as implemented in ADOL-C computes Gateaux derivatives in certain directions, which reduce to Fréchet derivatives only if the dependence on the direction is linear. Otherwise, the directional derivatives are merely positively homogeneous with respect to the scaling of the directions. For the reverse mode, ADOL-C sets the derivative of fabs(x) at the origin somewhat arbitrarily to zero.

We have defined binary functions fmin and fmax for adouble arguments, so that function and derivative values are obtained consistent with those of fabs according to the identities

$$\min(a,b) = [a+b-|a-b|]/2 \quad \text{and} \quad \max(a,b) = [a+b+|a-b|]/2 \quad .$$

These relations cannot hold if either a or b is infinite, in which case fmin or fmax and their derivatives may still be well defined. It should be noted that the directional differentiation of fmin and fmax yields at ties $a = b$ different results from the corresponding assignment based on the sign of $a - b$. For example, the statement

$$\mathsf{if\ (a < b)\ c = a;\ else\ c = b;}$$

yields for $\mathsf{a} = \mathsf{b}$ and $\mathsf{a}' < \mathsf{b}'$ the incorrect directional derivative value $\mathsf{c}' = \mathsf{b}'$ rather than the correct $\mathsf{c}' = \mathsf{a}'$. Therefore this form of conditional assignment should be avoided by use of the function fmin(a, b). There are also versions of fmin and fmax for two passive arguments and mixed passive/active arguments are handled by implicit conversion. On the function class obtained by composing the modulus with real analytic functions, the concept of directional differentiation can be extended to the propagation of unique one-sided Taylor expansions. The branches taken by fabs, fmin, and fmax, are recorded on the tape.

The functions sqrt, pow, and some inverse trigonometric functions have infinite slopes at the boundary points of their domains. At these marginal points the derivatives are set by ADOL-C to either ±InfVal, 0 or NoNum, where InfVal and NoNum are user-defined parameters, see Section 3.2. On IEEE machines InfVal can be set to the special value Inf = 1.0/0.0 and NoNum to NaN = 0.0/0.0. For example, at $\mathsf{a} = 0$ the first derivative b' of $\mathsf{b} = \mathsf{sqrt(a)}$ is set to

$$\mathsf{b}' = \begin{cases} \mathsf{InfVal} & \text{if } \mathsf{a}' > 0 \\ 0 & \text{if } \mathsf{a}' = 0 \\ \mathsf{NoNum} & \text{if } \mathsf{a}' < 0 \end{cases} \quad .$$

In other words, we consider a and consequently b as a constant when a' or more generally all computed Taylor coefficients are zero.

The general power function $\mathsf{pow(x,y)} = \mathsf{x}^{\mathsf{y}}$ is computed whenever it is defined for the corresponding double arguments. If x is negative, however, the partial derivative with respect

to an integral exponent is set to zero. The derivatives of the step functions floor, ceil, frexp, and ldexp are set to zero at all arguments x. The result values of the step functions are recorded on the tape and can later be checked to recognize whether a step to another level was taken during a forward sweep at different arguments than at taping time.

Some C implementations supply other special functions, in particular the error function erf(x). For the latter, we have included an adouble version in <adolc/adouble.cpp>, which has been commented out for systems on which the double valued version is not available. The increment and decrement operators ++, −− (prefix and postfix) are available for adoubles. Ambiguous statements like a += a++; must be avoided because the compiler may sequence the evaluation of the overloaded expression differently from the original in terms of doubles.

As we have indicated above, all subroutines called with active arguments must be modified or suitably overloaded. The simplest procedure is to declare the local variables of the function as active so that their internal calculations are also recorded on the tape. Unfortunately, this approach is likely to be unnecessarily inefficient and inaccurate if the original subroutine evaluates a special function that is defined as the solution of a particular mathematical problem. The most important examples are implicit functions, quadratures, and solutions of ordinary differential equations. Often the numerical methods for evaluating such special functions are elaborate, and their internal workings are not at all differentiable in the data. Rather than differentiating through such an adaptive procedure, one can obtain first and higher derivatives directly from the mathematical definition of the special function. Currently this direct approach has been implemented only for user-supplied quadratures as described in Section 8.5.

2.6 Reusing the Tape for Arbitrary Input Values

In some situations it may be desirable to calculate the value and derivatives of a function at arbitrary arguments by using a tape of the function evaluation at one argument and reevaluating the function and its derivatives using the given ADOL-C routines. This approach can significantly reduce run times, and it also allows to port problem functions, in the form of the corresponding tape files, into a computing environment that does not support C++ but does support C or Fortran. Therefore, the routines provided by ADOL-C for the evaluation of derivatives can be used to at arguments x other than the point at which the tape was generated, provided there are no user defined quadratures and all comparisons involving adoubles yield the same result. The last condition implies that the control flow is unaltered by the change of the independent variable values. Therefore, this sufficient condition is tested by ADOL-C and if it is not met the ADOL-C routine called for derivative calculations indicates this contingency through its return value. Currently, there are six return values, see Table 1.

In Figure 3 these return values are illustrated. If the user finds the return value of an ADOL-C routine to be negative the taping process simply has to be repeated by executing

+3	The function is locally analytic.
+2	The function is locally analytic but the sparsity structure (compared to the situation at the taping point) may have changed, e.g. while at taping arguments fmax(a,b) returned a we get b at the argument currently used.
+1	At least one of the functions fmin, fmax or fabs is evaluated at a tie or zero, respectively. Hence, the function to be differentiated is Lipschitz-continuous but possibly non-differentiable.
0	Some arithmetic comparison involving adoubles yields a tie. Hence, the function to be differentiated may be discontinuous.
-1	An adouble comparison yields different results from the evaluation point at which the tape was generated.
-2	The argument of a user-defined quadrature has changed from the evaluation point at which the tape was generated.

Table 1: Description of return values

the active section again. The crux of the problem lies in the fact that the tape records only the operations that are executed during one particular evaluation of the function. It also has no way to evaluate integrals since the corresponding quadratures are never recorded on the tape. Therefore, when there are user-defined quadratures the retaping is necessary at each new point. If there are only branches conditioned on adouble comparisons one may hope that re-taping becomes unnecessary when the points settle down in some small neighborhood, as one would expect for example in an iterative equation solver.

2.7 Conditional Assignments

It appears unsatisfactory that, for example, a simple table lookup of some physical property forces the re-recording of a possibly much larger calculation. However, the basic philosophy of ADOL-C is to overload arithmetic, rather than to generate a new program with jumps between "instructions", which would destroy the strictly sequential tape access and require the infusion of substantial compiler technology. Therefore, we introduce the two constructs of conditional assignments and active integers as partial remedies to the branching problem.

In many cases, the functionality of branches can be replaced by conditional assignments. For this purpose, we provide a special function called condassign(a,b,c,d). Its calling sequence corresponds to the syntax of the conditional assignment

$$\mathsf{a} = (\mathsf{b} > 0) \ ? \ \mathsf{c} \ : \ \mathsf{d};$$

which C++ inherited from C. However, here the arguments are restricted to be active or passive scalar arguments, and all expression arguments are evaluated before the test on b, which is different from the usual conditional assignment or the code segment.

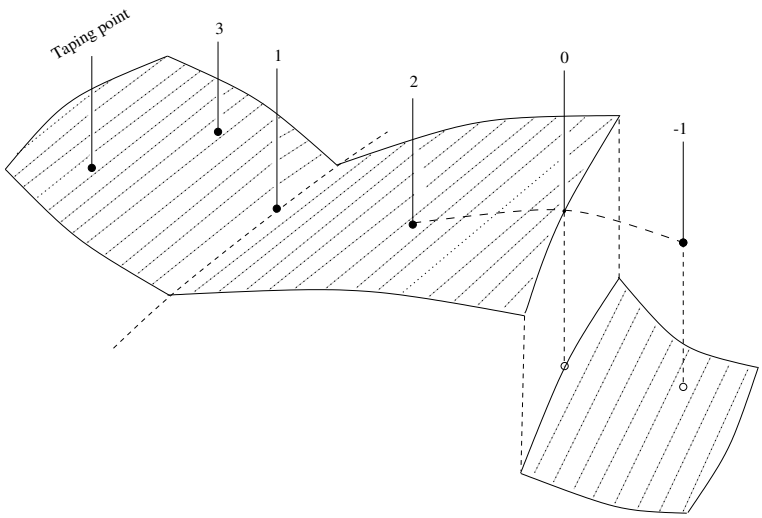

Figure 3: Return values around the taping point

Suppose the original program contains the code segment

$$\mathsf{if}\ (\mathsf{b} > 0)\ \mathsf{a} = \mathsf{c};\ \mathsf{else}\ \mathsf{a} = \mathsf{d};$$

Here, only one of the expressions (or, more generally, program blocks) c and d is evaluated, which exactly constitutes the problem for ADOL-C. To obtain the correct value a with ADOL-C, one may first execute both branches and then pick either c or d using condassign(a,b,c,d). To maintain consistency with the original code, one has to make sure that the two branches do not have any side effects that can interfere with each other or may be important for subsequent calculations. Furthermore the test parameter b has to be an adouble or an adouble expression. Otherwise the test condition b is recorded on the tape as a *constant* with its run time value. Thus the original dependency of b on active variables gets lost, for instance if b is a comparison expression, see Section 2.5. If there is no else part in a conditional assignment, one may call the three argument version condassign(a,b,c), which is logically equivalent to condassign(a,b,c,a) in that nothing happens if b is non-positive. The header file `<adolc/adouble.h>` contains also corresponding definitions of condassign(a,b,c,d) and condassign(a,b,c) for passive double arguments so that the modified code without any differentiation can be tested for correctness.

2.8 Step-by-Step Modification Procedure

To prepare a section of given C or C++ code for automatic differentiation as described above, one applies the following step-by-step procedure.

1. Use the statements trace_on(tag) or trace_on(tag,keep) and trace_off() or trace_off(file)

to mark the beginning and end of the active section.

2. Select the set of active variables, and change their type from double or float to adouble.
3. Select a sequence of independent variables, and initialize them with $\ll=$ assignments from passive variables or vectors.
4. Select a sequence of dependent variables among the active variables, and pass their final values to passive variable or vectors thereof by $\gg=$ assignments.
5. Compile the codes after including the header file `<adolc/adouble.h>`.

Typically, the first compilation will detect several type conflicts – usually attempts to convert from active to passive variables or to perform standard I/O of active variables. Since all standard C programs can be activated by a mechanical application of the procedure above, the following section is of importance only to advanced users.

3 Numbering the Tapes and Controlling the Buffer

The trace generated by the execution of an active section may stay within a triplet of internal arrays or it may be written out to corresponding files. We will refer to these triplets as the tape array or tape file, which may subsequently be used to evaluate the underlying function and its derivatives at the original point or at alternative arguments. If the active section involves user-defined quadratures it must be executed and re-taped at each new argument. Similarly, if conditions on adouble values lead to a different program branch being taken at a new argument the evaluation process also needs to be re-taped at the new point. Otherwise, direct evaluation from the tape by the routine function (Section 4.1) is likely to be faster. The use of quadratures and the results of all comparisons on adoubles are recorded on the tape so that function and other forward routines stop and return appropriate flags if their use without prior re-taping is unsafe. To avoid any re-taping certain types of branches can be recorded on the tape through the use of conditional assignments described before in Section 2.7.

Several tape files may be generated and kept simultaneously. A tape array is used as a triplet of buffers for a tape file if the length of any of the buffers exceeds the array length of BUFSIZE. This parameter is defined in the header file `<adolc/usrparms.h>` and may be adjusted by the user. For simple usage, trace_on may be called with only the tape tag as argument, and trace_off may be called without argument.

The optional integer argument keep of trace_on determines whether the numerical values of all active variables are recorded in a buffered temporary file called value stack before they will be overwritten. This option takes effect if keep = 1 and prepares the scene for an immediately following gradient evaluation by a call to a routine implementing the reverse mode as described in the Sections 5 and 6. The name of the temporary file is FNAME3

found in <adolc/dvlparms.h>. Alternatively, gradients may be evaluated by a call to gradient, which includes a preparatory forward sweep for the creation of the temporary file. If omitted, the argument keep defaults to 0, so that no temporary value stack file is generated.

By setting the optional integer argument file of trace_off to 1, the user may force a numbered tape file to be written even if the tape array (buffer) does not overflow. If the argument file is omitted, it defaults to 0, so that the tape array is written onto a tape file only if the length of any of the buffers exceeds BUFSIZE elements.

After the execution of an active section, if a tape file was generated, i.e., if the length of some buffer exceeded BUFSIZE elements or if the argument file of trace_off was set to 1, the files will be saved in the current working directory under the names FNAME.<tag>, FNAME1.<tag>, and FNAME2.<tag>, where tag is the mandatory argument to trace_on and FNAME, FNAME1, and FNAME2 are the tape file names found in <adolc/dvlparms.h>. Later, all problem-independent routines like gradient, jacobian, forward, reverse, and others expect as first argument a tag to determine the tape on which their respective computational task is to be performed. By calling trace_on with different tape tags, one can create several tapes for various function evaluations and subsequently perform function and derivative evaluations on one or more of them.

For example, suppose one wishes to calculate for two smooth functions $f_1(x)$ and $f_2(x)$

$$f(x) = \max\{f_1(x), f_2(x)\}, \qquad \nabla f(x),$$

and possibly higher derivatives where the two functions do not tie. Provided f_1 and f_2 are evaluated in two separate active sections, one can generate two different tapes by calling trace_on with tag = 1 and tag = 2 at the beginning of the respective active sections. Subsequently, one can decide whether $f(x) = f_1(x)$ or $f(x) = f_2(x)$ at the current argument and then evaluate the gradient $\nabla f(x)$ by calling gradient with the appropriate argument value tag = 1 or tag = 2.

3.1 Examining the Tape and Predicting Storage Requirements

At any point in the program, one may call the routine

```
void tapestats(unsigned short tag, int* counts)
```

with counts beeing an array of at least eleven integers. The first argument tag specifies the particular tape of interest. The components of counts represent

counts[0]:	the number of independents, i.e. calls to $\ll=$,
counts[1]:	the number of dependents, i.e. calls to $\gg=$,
counts[2]:	the maximal number of live active variables,
counts[3]:	the size of value stack (number of overwrites),
counts[4]:	the buffer size (a multiple of eight),

counts[5]: the total number of operations recorded,
counts[6-10]: other internal information about the tape.

The values maxlive = counts[2] and vssize = counts[3] determine the temporary storage requirements during calls to the routines implementing the forward and the reverse mode. For a certain degree deg ≥ 0, the scalar version of the forward mode involves apart from the tape buffers an array of (deg+1)∗maxlive doubles in core and, in addition, a sequential data set called the value stack of vssize∗keep revreals if called with the option keep > 0. Here the type revreal is defined as double or float in the header file `<adolc/usrparms.h>`. The latter choice halves the storage requirement for the sequential data set, which stays in core if its length is less than TBUFSIZE bytes and is otherwise written out to a temporary file. The parameter TBUFSIZE is defined in the header file `<adolc/usrparms.h>`. The drawback of the economical revreal = float choice is that subsequent calls to reverse mode implementations yield gradients and other adjoint vectors only in single-precision accuracy. This may be acceptable if the adjoint vectors represent rows of a Jacobian that is used for the calculation of Newton steps. In its scalar version, the reverse mode implementation involves the same number of doubles and twice as many revreals as the forward mode implementation. The storage requirements of the vector versions of the forward mode and reverse mode implementation are equal to that of the scalar versions multiplied by the vector length.

3.2 Customizing ADOL-C

Based on the information provided by the routine tapestats, the user may alter the following types and constant dimensions in the header file `<adolc/usrparms.h>` to suit his problem and environment.

BUFSIZE: This integer determines the length of internal buffers (default: 65 536). If the buffers are large enough to accommodate all required data, any file access is avoided unless trace_off is called with a positive argument. This desirable situation can be achieved for many problem functions with an execution trace of moderate size. Primarily BUFSIZE occurs as an argument to malloc, so that setting it unnecessarily large may have no ill effects, unless the operating system prohibits or penalizes large array allocations.

TBUFSIZE: This integer determines the length of the internal buffer for a value stack (default: 65 536).

TBUFNUM: This integer determines the maximal number of value stacks (default: 32).

locint: The range of the integer type locint determines how many adoubles can be simultaneously alive (default: unsigned short). In extreme cases when there are more than 65 536 adoubles alive at any one time, the type locint must be changed to unsigned int or unsigned long.

revreal: The choice of this floating-point type trades accuracy with storage for reverse sweeps (default: double). While functions and their derivatives are always evaluated in double precision during forward sweeps, gradients and other adjoint vectors are obtained with the precision determined by the type revreal. The less accurate choice revreal = float nearly halves the storage requirement during reverse sweeps.

fint: The integer data type used by Fortran callable versions of functions.

fdouble: The floating point data type used by Fortran callable versions of functions.

inf_num: This together with inf_den sets the "vertical" slope InfVal = inf_num/inf_den of special functions at the boundaries of their domains (default: inf_num = 1.0). On IEEE machines the default setting produces the standard Inf. On non-IEEE machines change these values to produce a small InfVal value and compare the results of two forward sweeps with different InfVal settings to detect a "vertical" slope.

inf_den: See inf_num (default: 0.0).

non_num: This together with non_den sets the mathematically undefined derivative value NoNum = non_num/non_den of special functions at the boundaries of their domains (default: non_num = 0.0). On IEEE machines the default setting produces the standard NaN. On non-IEEE machines change these values to produce a small NoNum value and compare the results of two forward sweeps with different NoNum settings to detect the occurrence of undefined derivative values.

non_den: See non_num (default: 0.0).

ADOLC_EPS: For testing on small numbers to avoid overflows (default: 10E-20).

ATRIG_ERF: By removing the comment signs the overloaded versions of the inverse hyperbolic functions and the error function are enabled (default: undefined).

DIAG_OUT: File identifier used as standard output for ADOL-C diagnostics (default: stdout).

ADOLC_USE_CALLOC: Selects the memory allocation routine used by ADOL-C. Malloc will be used if this variable is undefined. ADOLC_USE_CALLOC is defined by default to avoid incorrect result caused by uninitialized memory.

3.3 Warnings and Suggestions for Improved Efficiency

Since the type adouble has a nontrivial constructor, the mere declaration of large adouble arrays may take up considerable run time. The user should be warned against the usual Fortran practice of declaring fixed-size arrays that can accommodate the largest possible case of an evaluation program with variable dimensions. If such programs are converted to or written in C, the overloading in combination with ADOL-C will lead to very large run

time increases for comparatively small values of the problem dimension, because the actual computation is completely dominated by the construction of the large adouble arrays. The user is advised to create dynamic arrays of adoubles by using the C++ operator new and to destroy them using delete. For storage efficiency it is desirable that dynamic objects are created and destroyed in a last-in-first-out fashion.

Whenever an adouble is declared, the constructor for the type adouble assigns it a nominal address, which we will refer to as its *location*. The location is of the type locint defined in the header file <adolc/usrparms.h>. Active vectors occupy a range of contiguous locations. As long as the program execution never involves more than 65 536 active variables, the type locint may be defined as unsigned short. Otherwise, the range may be extended by defining locint as (unsigned) int or (unsigned) long, which may nearly double the overall mass storage requirement. Sometimes one can avoid exceeding the accessible range of unsigned shorts by using more local variables and deleting adoubles or adoublevs created by the new operator in a last-in-first-out fashion. When memory for adoubles is requested through a call to malloc() or other related C memory-allocating functions, the storage for these adoubles is allocated; however, the C++ adouble constructor is never called. The newly defined adoubles are never assigned a location and are not counted in the stack of live variables. Thus, any results depending upon these pseudo-adoubles will be incorrect. For these reasons **DO NOT use malloc() and related C memory-allocating functions when declaring adoubles (see the following paragraph).**

When an adouble goes out of scope or is explicitly deleted, the destructor notices that its location(s) may be freed for subsequent (nominal) reallocation. In general, this is not done immediately but is delayed until the locations to be deallocated form a contiguous tail of all locations currently being used.

As a consequence of this allocation scheme, the currently alive adouble locations always form a contiguous range of integers that grows and shrinks like a stack. Newly declared adoubles are placed on the top so that vectors of adoubles obtain a contiguous range of locations. While the C++ compiler can be expected to construct and destruct automatic variables in a last-in-first-out fashion, the user may upset this desirable pattern by deleting free-store adoubles too early or too late. Then the adouble stack may grow unnecessarily, but the numerical results will still be correct, unless an exception occurs because the range of locint is exceeded. In general, free-store adoubles should be deleted in a last-in-first-out fashion toward the end of the program block in which they were created. When this pattern is maintained, the maximum number of adoubles alive and, as a consequence, the randomly accessed storage space of the derivative evaluation routines is bounded by a small multiple of the memory used in the relevant section of the original program. Failure to delete dynamically allocated adoubles may cause that the maximal number of adoubles alive at one time will be exceeded if the same active section is called repeatedly. The same effect occurs if static adoubles are used.

To avoid the storage and manipulation of structurally trivial derivative values, one should pay careful attention to the naming of variables. Ideally, the intermediate values

generated during the evaluation of a vector function should be assigned to program variables that are consistently either active or passive, in that all their values either are or are not dependent on the independent variables in a nontrivial way. For example, this rule is violated if a temporary variable is successively used to accumulate inner products involving first only passive and later active arrays. Then the first inner product and all its successors in the data dependency graph become artificially active and the derivative evaluation routines described later will waste time allocating and propagating trivial or useless derivatives. Sometimes even values that do depend on the independent variables may be of only transitory importance and may not affect the dependent variables. For example, this is true for multipliers that are used to scale linear equations, but whose values do not influence the dependent variables in a mathematical sense. Such dead-end variables can be deactivated by the use of the value function, which converts adoubles to doubles. The deleterious effects of unnecessary activity are partly alleviated by run time activity flags in the derivative routine hov_reverse presented in Section 5.

4 Easy-To-Use Drivers

For the convenience of the user, ADOL-C provides several easy-to-use drivers that compute the most frequently required derivative objects. Throughout, we assume that after the execution of an active section, the corresponding tape with the identifier tag contains a detailed record of the computational process by which the final values y of the dependent variables were obtained from the values x of the independent variables. We will denote this functional relation between the input variables x and the output variables y by

$$F : \mathbb{R}^n \mapsto \mathbb{R}^m, \qquad x \to F(x) \equiv y.$$

The return value of all drivers presented in this section indicate the validity of the tape as explained in Section 2.6. The presented drivers are all C functions and therefore can be used within C and C++ programs. Some Fortran-callable companions can be found in the appropriate header files.

4.1 Drivers for Optimization and Nonlinear Equations

The drivers provided for solving optimization problems and nonlinear equations are prototyped in the header file `<adolc/drivers/drivers.h>`, which is included automatically by the global header file `<adolc/adolc.h>` (see Section 8.3).

The routine function allows to evaluate the desired function from the tape instead of executing the corresponding source code:

```
int function(tag,m,n,x,y)
short int tag;                  // tape identification
```

```
int m;                  // number of dependent variables m
int n;                  // number of independent variables n
double x[n];            // independent vector x
double y[m];            // dependent vector y = F(x)
```

If the original evaluation program is available this double version should be used to compute the function value in order to avoid the interpretative overhead.

For the calculation of whole derivative vectors and matrices up to order 2 there are the following procedures:

```
int gradient(tag,n,x,g)
short int tag;          // tape identification
int n;                  // number of independent variables n and m = 1
double x[n];            // independent vector x
double g[n];            // resulting gradient ∇F(x)

int jacobian(tag,m,n,x,J)
short int tag;          // tape identification
int m;                  // number of dependent variables m
int n;                  // number of independent variables n
double x[n];            // independent vector x
double J[m][n];         // resulting Jacobian F'(x)

int hessian(tag,n,x,H)
short int tag;          // tape identification
int n;                  // number of independent variables n and m = 1
double x[n];            // independent vector x
double H[n][n];         // resulting Hessian matrix ∇²F(x)
```

The driver routine hessian computes only the lower half of $\nabla^2 f(x_0)$ so that all values H[i][j] with $j > i$ of H allocated as a square array remain untouched during the call of hessian. Hence only $i + 1$ doubles need to be allocated starting at the position H[i].

To use the full capability of automatic differentiation when the product of derivatives with certain weight vectors or directions are needed, ADOL-C offers the following four drivers:

```
int vec_jac(tag,m,n,repeat,x,u,z)
short int tag;          // tape identification
int m;                  // number of dependent variables m
int n;                  // number of independent variables n
int repeat;             // indicate repeated call at same argument
```

```
double x[n];          // independent vector x
double u[m];          // range weight vector u
double z[n];          // result z = u^T F'(x)
```

If a nonzero value of the parameter repeat indicates that the routine vec_jac has been called at the same argument immediately before, the internal forward mode evaluation will be skipped and only reverse mode evaluation with the corresponding arguments is executed resulting in a reduced computational complexity of the function vec_jac.

```
int jac_vec(tag,m,n,x,v,z)
short int tag;        // tape identification
int m;                // number of dependent variables m
int n;                // number of independent variables n
double x[n];          // independent vector x
double v[n];          // tangent vector v
double z[m];          // result z = F'(x)v

int hess_vec(tag,n,x,v,z)
short int tag;        // tape identification
int n;                // number of independent variables n
double x[n];          // independent vector x
double v[n];          // tangent vector v
double z[n];          // result z = ∇²F(x)v

int hess_mat(tag,n,p,x,V,Z)
short int tag;        // tape identification
int n;                // number of independent variables n
int p;                // number of columns in V
double x[n];          // independent vector x
double V[n][p];       // tangent matrix V
double Z[n][p];       // result Z = ∇²F(x)V

int lagra_hess_vec(tag,m,n,x,v,u,h)
short int tag;        // tape identification
int m;                // number of dependent variables m
int n;                // number of independent variables n
double x[n];          // independent vector x
double v[n];          // tangent vector v
double u[m];          // range weight vector u
double h[n];          // result h = u^T ∇²F(x)v
```

The next procedure allows the user to perform Newton steps only having the corresponding tape at hand:

```
int jac_solv(tag,n,x,b,sparse,mode)
short int tag;          // tape identification
int n;                  // number of independent variables n
double x[n];            // independent vector x as
double b[n];            // in: right-hand side b, out: result w of F(x)w = b
int sparse;             // option to use sparsity
int mode;               // option to choose different solvers
```

On entry, parameter b of the routine jac_solv contains the right-hand side of the equation $F(x)w = b$ to be solved. On exit, b equals the solution w of this equation. Currently, the parameter sparse is not used, but will be subject of further developments of ADOL-C in the future. It is planed to provide a sparse version of jac_solv. If mode $= 0$ only the Jacobian of the function given by the tape labeled with tag is provided internally. The LU-factorization of this Jacobian is computed for mode $= 1$. The solution of the equation is calculated if mode $= 2$. Hence, it is possible to compute the LU-factorization only once. Then the equation can be solved for several right-hand sides b without calculating the Jacobian and its factorization again.

If the original evaluation code of a function contains neither quadratures nor branches, all drivers described above can be used to evaluate derivatives at any argument in its domain. The same still applies if there are no user defined quadratures and all comparisons involving adoubles have the same result as during taping. If this assumption is falsely made all drivers while internally calling the forward mode evaluation will return the value -1 or -2 as already specified in Section 2.6.

4.2 Drivers for Ordinary Differential Equations

When F is the right-hand side of an (autonomous) ordinary differential equation

$$x'(t) \;=\; F(x(t)),$$

we must have $m = n$. Along any solution path $x(t)$ its Taylor coefficients x_j at some time, e.g., $t = 0$, must satisfy the relation

$$x_{i+1} = \frac{1}{1+i} y_i.$$

with the y_j the Taylor coefficients of its derivative $y(t) = x'(t)$, namely,

$$y(t) \;\equiv\; F(x(t)) \;:\; I\!R \;\mapsto\; I\!R^m$$

defined by an autonomous right-hand side F recorded on the tape. Using this relation, one can generate the Taylor coefficients x_i, $i \leq deg$, recursively from the current point x_0. This task is achieved by the driver routine forode defined as follows:

```
int forode(tag,n,tau,dol,deg,X)
short int tag;          // tape identification
int n;                  // number of state variables n
double tau;             // scaling parameter
int dol;                // degree on previous call
int deg;                // degree on current call
double X[n][deg+1];     // Taylor coefficient vector X
```

If dol is positive, it is assumed that forode has been called before at the same point so that all Taylor coefficient vectors up to the dol-th are already correct.

Subsequently one may call the driver routine reverse or corresponding low level routines as explained in the Sections 6 and 5, respectively, to compute the family of square matrices Z[n][n][deg] defined by

$$Z_j \equiv U \frac{\partial y_j}{\partial x_0} \in I\!R^{q \times n},$$

with double** U= I_n the identity matrix of order n.

For the numerical solutions of ordinary differential equations, one may also wish to calculate the Jacobians

$$B_j \;\equiv\; \frac{\mathrm{d}x_{j+1}}{\mathrm{d}x_0} \;\in\; I\!R^{n \times n}\,, \tag{1}$$

which exist provided F is sufficiently smooth. These matrices can be obtained from the partial derivatives $\partial y_i / \partial x_0$ by an appropriate version of the chain rule. To compute the total derivatives $B = (B_j)_{0 \le j < d}$ defined in (1), one has to evaluate $\frac{1}{2}d(d-1)$ matrix-matrix products. This can be done by a call of the routine accode after the corresponding evaluation of the hov_reverse function. The interface of accode is defined as follows:

```
int accode(n,tau,deg,Z,B,nz)
int n;                  // number of state variables n
double tau;             // scaling parameter
int deg;                // degree on current call
double Z[n][n][deg];    // partials of coefficient vectors
double B[n][n][deg];    // result B as defined in (1)
short nz[n][n];         // optional nonzero pattern
```

Sparsity information can be exploited by accode using the array nz. For this purpose, nz has to be set by a call of the routine reverse or the corresponding basic routines as explained below in Section 5 and Section 6, respectively. The non-positive entries of nz are then changed by accode so that upon return

$$\mathsf{B[i][j][k]} \;\equiv\; 0 \quad \text{if} \quad \mathsf{k} \le -\mathsf{nz[i][j]}\;.$$

In other words, the matrices $B_k = \mathsf{B[\,][\,][k]}$ have a sparsity pattern that fills in as k grows. Note, that there need to be no loss in computational efficiency if a time-dependent ordinary differential equation is rewritten in autonomous form.

The prototype of the ODE-drivers forode and accode is contained in the header file `<adolc/drivers/odedrivers.h>`. The global header file `<adolc/adolc.h>` includes this file automatically, see Section 8.3.

An example program using the procedures forode and accode together with more detailed information about the coding can be found in Section 9.4. The corresponding source code `odexam.cpp` is contained in the subdirectory `examples`.

4.3 Drivers for Sparse Jacobians and Sparse Hessians

Quite often, the Jacobians and Hessians that have to be computed are sparse matrices. Therefore, ADOL-C provides additionally drivers that allow the exploitation of sparsity.

Sparse Jacobians and Sparse Hessians

To compute the entries of sparse Jacobians and sparse Hessians, respectively, in coordinate format one may use the drivers:

```
int sparse_jac(tag,m,n,repeat,x,&nnz,&rind,&cind,&values)
short int tag;                  // tape identification
int m;                          // number of dependent variables m
int n;                          // number of independent variables n
int repeat;                     // indicate repeated call at same argument
double x[n];                    // independent vector x
int nnz;                        // number of nonzeros
unsigned int rind[nnz];         // row index
unsigned int cind[nnz];         // column index
double values[nnz];             // non-zero values

int sparse_hess(tag,n,repeat,x,&nnz,&rind,&cind,&values)
short int tag;                  // tape identification
int n;                          // number of independent variables n and m = 1
int repeat;                     // indicate repeated call at same argument
double x[n];                    // independent vector x
int nnz;                        // number of nonzeros
unsigned int rind[nnz];         // row indices
unsigned int cind[nnz];         // column indices
double values[nnz];             // non-zero values
```

Once more, the input variables are the identifier for the internal representation tag, if required the number of dependents m, and the number of independents n for a consistency check. Furthermore, the flag repeat=0 indicates that the functions are called at a point with

a new sparsity structure, whereas repeat=1 results in the re-usage of the sparsity pattern from the previous call. The current values of the independents are given by the array x. The input/output variable nnz stores the number of the nonzero entries. Therefore, nnz denotes also the length of the arrays r_ind storing the row indices, c_ind storing the column indices, and values storing the values of the nonzero entries. If sparse_jac and sparse_hess are called with repeat=0, the functions determine the number of nonzeros for the sparsity pattern defined by the value of x, allocate appropriate arrays r_ind, c_ind, and values and store the desired information in these arrays. During the next function call with repeat=1 the allocated memory is reused such that only the values of the arrays are changed. Before calling sparse_jac or sparse_hess once more with repeat=0 the user is responsible for the deallocation of the array r_ind, c_ind, and values using the function free()!

This driver routine is prototyped in the header file `<adolc/sparse/sparse.h>`, which is included automatically by the global header file `<adolc/adolc.h>` (see Section 8.3). An example code illustrating the usage of sparse_jac and sparse_hess can be found in the file `sparse_jac_hess_exam.cpp` contained in `examples/additional_examples/sparse`.

Computation of Sparsity Pattern

ADOL-C offers a convenient way of determining the sparsity structure of a Jacobian matrix using the function:

```
int jac_pat(tag, m, n, x, rb, cb, JP, options)
short int tag;                  // tape identification
int m;                          // number of dependent variables m
int n;                          // number of independent variables n
double x[n];                    // independent variables x_0
unsigned int rb[];              // description of dependents
unsigned int cb[];              // description of independents
unsigned int JP[][];            // row compressed sparsity structure
int options[2];                 // array of control parameters
```

If rb = NULL and cb = NULL, the sparsity pattern of the Jacobian is computed in a compressed row format. For this purpose, JP has to be an m dimensional array of pointers to unsigned ints, i.e., one has unsigned int* JP[m]. During the call of jac_pat, the number $\hat{n}_i$ of nonzero entries in row i of the Jacobian is determined for all $1 \leq i \leq m$. Then, a memory allocation is performed such that JP[i-1] points to a block of $\hat{n}_i + 1$ unsigned int for all $1 \leq i \leq m$ and JP[i-1][0] is set to $\hat{n}_i$. Subsequently, the column indicies of the j nonzero entries in the ith row are stored in the components JP[i-1][1], ..., JP[i-1][j].

The elements of the array options control the action of jac_pat according to Table 2. The value of options[0] selects the direction of bit pattern propagation. Depending on the number of independent n and of dependent variables m one would prefer the forward mode

options[0]	0	mode selection (default)
	1	forward mode
	2	reverse mode
options[1]	0	safe mode (default)
	1	tight mode

Table 2: jac_pat parameter options

if n is significant smaller than m and would otherwise use the reverse mode. The component options[1] determines the usage of the safe or tight mode of bit pattern propagation. The first, more conservative option is the default. It accounts for all dependences that might occur for any value of the independent variables. For example, the intermediate c = max(a,b) is always assumed to depend on all independent variables that a or b dependent on, i.e. the bit pattern associated with c is set to the logical OR of those associated with a and b. In contrast the tight option gives this result only in the unlikely event of an exact tie a = b. Otherwise it sets the bit pattern associated with c either to that of a or to that of b, depending on whether c = a or c = b locally. Obviously, the sparsity pattern obtained with the tight option may contain more zeros than that obtained with the safe option. On the other hand, it will only be valid at points belonging to an area where the function F is locally analytic and that contains the point at which the internal representation was generated. Since generating the sparsity structure using the safe version does not require any reevaluation, it may thus reduce the overall computational cost despite the fact that it produces more nonzero entries.

For several applications, a grouping of the variables into blocks is quite obvious. For example, variable grouping happens in simulations that involve the physical space, where the three components of position or velocity vectors naturally occur together or not at all in various relations. Taking the group structure into account, the overall cost of computing the sparsity structure can be reduced using the arrays rb and cb. Here it is also possible to set one of them to NULL treating each variable as a singleton block. If not equal to NULL, the component rb[0] determines the number of blocks of dependent variables and cb[0] the number of blocks of independent variables. The number of the column group that contains the variable x[j] is represented by the integer cb[j+1] < cb[0]. Correspondingly, rb[i+1] < rb[0] represents the number of the row group that contains the variable y[i]. The derivative of y[i] with respect to x[j] may be nonzero only if the block with index (rb[i],cb[j]) as a whole is nonzero. Hence, the array JP contains the compressed sparsity structure in a format similar to the one explained above, where each row corresponds to a block of dependent variables. JP[][0] is always the number of blocks of independent variables on which the current block of dependent variables depends. The components JP[][i], $0 <$ i $\leq$ JP[][0] store the indices of these blocks of independent variables (see Fig. 4).

The routine jac_pat uses the propagation of bitpattern to determine the sparsity pattern. Therefore, a kind of "strip-mining" is used to cope with large matrix dimensions.

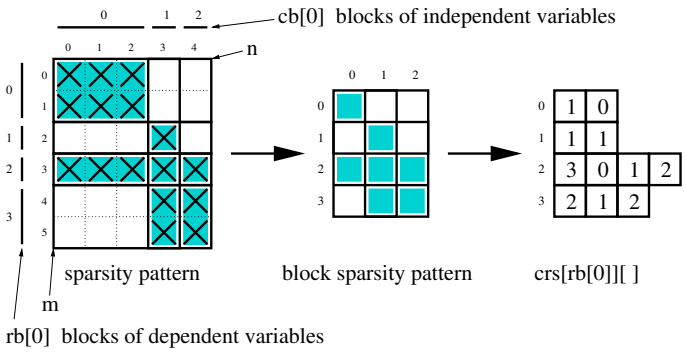

Figure 4: Jacobian (block) sparsity pattern

If the system happens to run out of memory, one may reduce the value of the constant PQ_STRIPMINE_MAX following the instructions in `<adolc/sparse/jacutils.h>`.

The driver routine is prototyped in the header file `<adolc/sparse/sparse.h>`, which is included automatically by the global header file `<adolc/adolc.h>` (see Section 8.3). The determination of sparsity patterns is illustrated by the code `sparse_jac_hess_exam.cpp` contained in `examples/additional_examples/sparse`. The example `jacpatexam.cpp` located in the same directory demonstrates the grouping of independents and dependents for the driver routine jac_pat

To compute the sparsity pattern of a Hessian in a row compressed form, ADOL-C provides the driver

```
int hess_pat(tag, n, x, HP, options)
short int tag;                  // tape identification
int n;                          // number of independent variables n
double x[n];                    // independent variables x_0
unsigned int HP[][];            // row compressed sparsity structure
int options[2];                 // array of control parameters
```

where the user has to provide HP as an n dimensional array of pointers to unsigned ints. After the function call HP contains the sparsity pattern, where HP[j][0] contains the number of nonzero elements in the jth row for $1 \leq j \leq n$. The components P[j][i], $0 <$i $\leq$ P[j][0] store the indices of these entries. For determining the sparsity pattern, ADOL-C uses the algorithm described in [25].

This driver routine is prototyped in the header file `<adolc/sparse/sparse.h>`, which is included automatically by the global header file `<adolc/adolc.h>` (see Section 8.3). The example `sparse_jac_hess_exam.cpp` employing the procedure hess_pat can be found in the file `sparse_jac_hess_exam.cpp` contained in `examples/additional_examples/sparse`.

Calculation of Seed Matrices

To compute a compressed derivative matrix from a given sparsity pattern, one has to calculate an appropriate seed matrix that can be used as input for the derivative calculation. To facilitate the generation of seed matrices for a sparsity pattern given in row compressed form, ADOL-C provides the following two drivers:

```
int generate_seed_jac(m, n, JP, S, p)
int m;                          // number of dependent variables m
int n;                          // number of independent variables n
unsigned int JP[][];            // row compressed sparsity structure of Jacobian
double S[n][p];                 // seed matrix
int p;                          // number of columns in S
```

The input variables are the number of dependent variables m, the number of independent variables n and the sparsity pattern JP of the Jacobian computed for example by jac_pat. First, generate_seed_jac performs a coloring of the bipartite graph defined by the sparsity pattern JP as described in [9]. The number of colors needed for the coloring determines the number of columns p in the seed matrix. Subsequently, generate_seed_jac allocates the memory needed by S and initializes S according to the graph coloring. The coloring algorithm that is applied in generate_seed_jac is used also by the driver sparse_jac described above.

```
int generate_seed_hess(n, HP, S, p)
int n;                          // number of independent variables n
unsigned int HP[][];            // row compressed sparsity structure of Jacobain
double S[n][p];                 // seed matrix
int p;                          // number of columns in S
```

The input variables are the number of independents n and the sparsity pattern HP of the Hessian computed for example by hess_pat. First, generate_seed_hess performs a coloring of the adjacent graph defined by the sparsity pattern HP as described in [9]. The number of colors needed for the coloring determines the number of columns p in the seed matrix. Subsequently, generate_seed_hess allocates the memory needed by S and initializes S according to the graph coloring. The coloring algorithm applied in generate_seed_hess is used also by the driver sparse_hess described above.

These driver routines are prototyped in the header file `<adolc/sparse/sparse.h>`, which is included automatically by the global header file `<adolc/adolc.h>` (see Section 8.3). An example code illustrating the usage of the procedures generate_seed_jac and generate_seed_hess can be found in the file `sparse_jac_hess_exam.cpp` contained in the subdirectory `examples/additional_examples/sparse`.

4.4 Higher Derivative Tensors

Many applications in scientific computing need second- and higher-order derivatives. Often, one does not require full derivative tensors but only the derivatives in certain directions $s_i \in \mathbb{R}^n$. Suppose a collection of p directions $s_i \in \mathbb{R}^n$ is given, which form a matrix

$$S \;=\; [s_1, s_2, \ldots, s_p] \;\in\; \mathbb{R}^{n \times p}.$$

One possible choice is $S = I_n$ with $p = n$, which leads to full tensors being evaluated. ADOL-C provides the function tensor_eval to calculate the derivative tensors

$$\nabla_S^k F(x_0) \;=\; \frac{\partial^k}{\partial z^k} F(x_0 + Sz)\bigg|_{z=0} \in \mathbb{R}^{p^k} \quad \text{for} \quad k = 0, \ldots, d \tag{2}$$

simultaneously. The function tensor_eval has the following calling sequence and parameters:

```
void tensor_eval(tag,m,n,d,p,x,tensor,S)
short int tag;              // tape identification
int m;                      // number of dependent variables m
int n;                      // number of independent variables n
int d;                      // highest derivative degree d
int p;                      // number of directions p
double x[n];                // values of independent variables x_0
double tensor[m][size];     // result as defined in (2) in compressed form
double S[n][p];             // seed matrix S
```

Using the symmetry of the tensors defined by (2), the memory requirement can be reduced enormously. The collection of tensors up to order d comprises $\binom{p+d}{d}$ distinct elements. Hence, the second dimension of tensor must be greater or equal to $\binom{p+d}{d}$. To compute the derivatives, tensor_eval propagates internally univariate Taylor series along $\binom{n+d-1}{d}$ directions. Then the desired values are interpolated. This approach is described in [14].

The access of individual entries in symmetric tensors of higher order is a little tricky. We always store the derivative values in the two dimensional array tensor and provide two different ways of accessing them. The leading dimension of the tensor array ranges over the component index i of the function F, i.e., F_{i+1} for $i = 0, \ldots, m-1$. The sub-arrays pointed to by tensor[i] have identical structure for all i. Each of them represents the symmetric tensors up to order d of the scalar function F_{i+1} in p variables. The $\binom{p+d}{d}$ mixed partial derivatives in each of the m tensors are linearly ordered according to the tetrahedral scheme described by Knuth [18]. In the familiar quadratic case $d = 2$ the derivative with respect to z_j and z_k with z as in (2) and $j \le k$ is stored at tensor[i][l] with $l = k * (k+1)/2 + j$. At $j = 0 = k$ and hence $l = 0$ we find the function value F_{i+1} itself and the gradient $\nabla F_{i+1} = \partial F_{i+1}/\partial x_k$ is stored at $l = k(k+1)/2$ with $j = 0$ for $k = 1, \ldots, p$.

For general d we combine the variable indices to a multi-index $j = (j_1, j_2, \ldots, j_d)$, where j_k indicates differentiation with respect to variable x_{j_k} with $j_k \in \{0, 1, \ldots, p\}$. The value

$j_k = 0$ indicates no differentiation so that all lower derivatives are also contained in the same data structure as described above for the quadratic case. The location of the partial derivative specified by j is computed by the function

```
int address(d, j)
int d;                    // highest derivative degree d
int j[d];                 // multi-index j
```

and it may thus be referenced as tensor[i][address(d, j)]. Notice that the address computation does depend on the degree d but not on the number of directions p, which could theoretically be enlarged without the need to reallocate the original tensor. Also, the components of j need to be non-increasing. To some C programmers it may appear more natural to access tensor entries by successive dereferencing in the form tensorentry[i][j1][j2]...[jd]. We have also provided this mode, albeit with the restriction that the indices $j_1, j_2, \ldots, j_d$ are non-increasing. In the second order case this means that the Hessian entries must be specified in or below the diagonal. If this restriction is violated the values are almost certain to be wrong and array bounds may be violated. We emphasize that subscripting is not overloaded but that tensorentry is a conventional and thus moderately efficient C pointer structure. Such a pointer structure can be allocated and set up completely by the function

```
void** tensorsetup(m,p,d,tensor)
int m;                    // number of dependent variables n
int p;                    // number of directions p
int d;                    // highest derivative degree d
double tensor[m][size];   // pointer to two dimensional array
```

Here, tensor is the array of m pointers pointing to arrays of size $\geq \binom{p+d}{d}$ allocated by the user before. During the execution of tensorsetup, $d-1$ layers of pointers are set up so that the return value allows the direct dereferencing of individual tensor elements.

For example, suppose some active section involving $m \geq 5$ dependents and $n \geq 2$ independents has been executed and taped. We may select $p = 2$, $d = 3$ and initialize the $n \times 2$ seed matrix S with two columns s_1 and s_2. Then we are able to execute the code segment

```
double**** tensorentry = (double****) tensorsetup(m,p,d,tensor);
tensor_eval(tag,m,n,d,p,x,tensor,S);
```

This way, we evaluated all tensors defined in (2) up to degree 3 in both directions s_1 and s_2 at some argument x. To allow the access of tensor entries by dereferencing the pointer structure tensorentry has been created. Now, the value of the mixed partial

$$\left.\frac{\partial^3 F_5(x + s_1 z_1 + s_2 z_2)}{\partial z_1^2 \partial z_2}\right|_{z_1=0=z_2}$$

can be recovered as

tensorentry[4][2][1][1] or tensor[4][address(d, j)],

where the integer array j may equal (1,1,2), (1,2,1) or (2,1,1). Analogously, the entry

tensorentry[2][1][0][0] or tensor[2][address(d, j)]

with j = (1,0,0) contains the first derivative of the third dependent variable F_3 with respect to the first differentiation parameter z_1.

Note, that the pointer structure tensorentry has to be set up only once. Changing the values of the array tensor, e.g. by a further call of tensor_eval, directly effects the values accessed by tensorentry. When no more derivative evaluations are desired the pointer structure tensorentry can be deallocated by a call to the function

```
int freetensor(m,p,d, (double ****) tensorentry)
int m;                          // number of dependent variables m
int p;                          // number of independent variables p
int d;                          // highest derivative degree d
double*** tensorentry[m];       // return value of tensorsetup
```

that does not deallocate the array tensor.

The drivers provided for efficient calculation of higher order derivatives are prototyped in the header file `<adolc/drivers/taylor.h>`, which is included by the global header file `<adolc/adolc.h>` automatically (see Section 8.3). Example codes using the above procedures can be found in the files `taylorexam.C` and `accessexam.C` contained in the subdirectory `examples/additional_examples/taylor`.

4.5 Derivatives of Implicit and Inverse Functions

Frequently, one needs derivatives of variables $y \in \mathbb{R}^m$ that are implicitly defined as functions of some variables $x \in \mathbb{R}^{n-m}$ by an algebraic system of equations

$$G(z) \;=\; 0 \in \mathbb{R}^m \quad \text{with} \quad z = (y, x) \in \mathbb{R}^n.$$

Naturally, the n arguments of G need not be partitioned in this regular fashion and we wish to provide flexibility for a convenient selection of the $n - m$ *truly* independent variables. Let $P \in \mathbb{R}^{(n-m)\times n}$ be a $0-1$ matrix that picks out these variables so that it is a column permutation of the matrix $[0, I_{n-m}] \in \mathbb{R}^{(n-m)\times n}$. Then the nonlinear system

$$G(z) \;=\; 0, \quad Pz = x,$$

has a regular Jacobian, wherever the implicit function theorem yields y as a function of x. Hence, we may also write

$$F(z) = \begin{pmatrix} G(z) \\ Pz \end{pmatrix} \equiv \begin{pmatrix} 0 \\ Pz \end{pmatrix} \equiv S\,x, \tag{3}$$

where $S = [0, I_p]^T \in \mathbb{R}^{n \times p}$ with $p = n - m$. Now, we have rewritten the original implicit functional relation between x and y as an inverse relation $F(z) = Sx$. In practice, we may implement the projection P simply by marking $n - m$ of the independents also dependent.

Given any $F : \mathbb{R}^n \mapsto \mathbb{R}^n$ that is locally invertible and an arbitrary seed matrix $S \in \mathbb{R}^{n \times p}$ we may evaluate all derivatives of $z \in \mathbb{R}^n$ with respect to $x \in \mathbb{R}^p$ by calling the following routine:

```
void inverse_tensor_eval(tag,n,d,p,z,tensor,S)
short int tag;              // tape identification
int n;                      // number of variables n
int d;                      // highest derivative degree d
int p;                      // number of directions p
double z[n];                // values of independent variables z
double tensor[n][size];     // partials of z with respect to x
double S[n][p];             // seed matrix S
```

The results obtained in tensor are exactly the same as if we had called tensor_eval with tag pointing to a tape for the evaluation of the inverse function $z = F^{-1}(y)$ for which naturally $n = m$. Note that the columns of S belong to the domain of that function. Individual derivative components can be accessed in tensor exactly as in the explicit case described above.

It must be understood that inverse_tensor_eval actually computes the derivatives of z with respect to x that is defined by the equation $F(z) = F(z_0) + S\,x$. In other words the base point at which the inverse function is differentiated is given by $F(z_0)$. The routine has no capability for inverting F itself as solving systems of nonlinear equations $F(z) = 0$ in the first place is not just a differentiation task. However, the routine jac_solv described in Section 4.1 may certainly be very useful for that purpose.

As an example consider the following two nonlinear expressions

$$\begin{aligned} G_1(z_1, z_2, z_3, z_4) &= z_1^2 + z_2^2 - z_3^2 \\ G_2(z_1, z_2, z_3, z_4) &= \cos(z_4) - z_1/z_3 \ . \end{aligned}$$

The equations $G(z) = 0$ describe the relation between the Cartesian coordinates (z_1, z_2) and the polar coordinates (z_3, z_4) in the plane. Now, suppose we are interested in the derivatives of the second Cartesian $y_1 = z_2$ and the second (angular) polar coordinate $y_2 = z_4$ with respect to the other two variables $x_1 = z_1$ and $x_2 = z_3$. Then the active section could look simply like

```
for (j=1; j < 5; j++)    z[j] <<= zp[j];
g[1] = z[1]*z[1]+z[2]*z[2]-z[3]*z[3];
g[2] = cos(z[4]) - z[1]/z[3];
g[1] >>= gp[1];          g[2] >>= gp[2];
z[1] >>= zd[1];          z[3] >>= zd[2];
```

where zd[1] and zd[2] are dummy arguments. In the last line the two independent variables z[1] and z[3] are made simultaneously dependent thus generating a square system that can be inverted (at most arguments). The corresponding projection and seed matrix are

$$P = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad \text{and} \quad S^T = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} .$$

Provided the vector zp is consistent in that its Cartesian and polar components describe the same point in the plane the resulting tuple gp must vanish. The call to inverse_tensor_eval with $n = 4$, $p = 2$ and d as desired will yield the implicit derivatives, provided tensor has been allocated appropriately of course and S has the value given above. The example is untypical in that the implicit function could also be obtained explicitly by symbolic manipulations. It is typical in that the subset of z components that are to be considered as truly independent can be selected and altered with next to no effort at all.

The presented drivers are prototyped in the header file `<adolc/drivers/taylor.h>`. As indicated before this header is included by the global header file `<adolc/adolc.h>` automatically (see Section 8.3). The example programs `inversexam.cpp`, `coordinates.cpp` and `trigger.cpp` in the directory `examples/additional_examples/taylor` show the application of the procedures described here.

5 Basic Drivers for the Forward and Reverse Mode

In this section, we present tailored drivers for different variants of the forward mode and the reverse mode, respectively. For a better understanding, we start with a short description of the mathematical background.

Provided no arithmetic exception occurs, no comparison including fmax or fmin yields a tie, fabs does not yield zero, and all special functions were evaluated in the interior of their domains, the functional relation between the input variables x and the output variables y denoted by $y = F(x)$ is in fact analytic. In other words, we can compute arbitrarily high derivatives of the vector function $F : I\!R^n \mapsto I\!R^m$ defined by the active section. We find it most convenient to describe and compute derivatives in terms of univariate Taylor expansions, which are truncated after the highest derivative degree d that is desired by the user. Let

$$x(t) \equiv \sum_{j=0}^{d} x_j t^j \; : \; I\!R \mapsto I\!R^n \tag{4}$$

denote any vector polynomial in the scalar variable $t \in I\!R$. In other words, $x(t)$ describes a path in $I\!R^n$ parameterized by t. The Taylor coefficient vectors

$$x_j \;=\; \frac{1}{j!}\,\frac{\partial^j}{\partial t^j} x(t)\bigg|_{t=0}$$

are simply the scaled derivatives of $x(t)$ at the parameter origin $t = 0$. The first two vectors $x_1, x_2 \in I\!R^n$ can be visualized as tangent and curvature at the base point x_0, respectively. Provided that F is d times continuously differentiable, it follows from the chain rule that the image path

$$y(t) \;\equiv\; F(x(t)) \;:\; I\!R \;\mapsto\; I\!R^m \tag{5}$$

is also smooth and has $(d+1)$ Taylor coefficient vectors $y_j \in I\!R^m$ at $t = 0$, so that

$$y(t) \;=\; \sum_{j=0}^{d} y_j t^j + O(t^{d+1}). \tag{6}$$

Also as a consequence of the chain rule, one can observe that each y_j is uniquely and smoothly determined by the coefficient vectors x_i with $i \le j$. In particular we have

$$\begin{aligned}
y_0 &= F(x_0) \\
y_1 &= F'(x_0)x_1 \\
y_2 &= F'(x_0)x_2 + \frac{1}{2}F''(x_0)x_1x_1 \\
y_3 &= F'(x_0)x_3 + F''(x_0)x_1x_2 + \frac{1}{6}F'''(x_0)x_1x_1x_1 \\
&\ldots
\end{aligned} \tag{7}$$

In writing down the last equations we have already departed from the usual matrix-vector notation. It is well known that the number of terms that occur in these "symbolic" expressions for the y_j in terms of the first j derivative tensors of F and the "input" coefficients x_i with $i \le j$ grows very rapidly with j. Fortunately, this exponential growth does not occur in automatic differentiation, where the many terms are somehow implicitly combined so that storage and operations count grow only quadratically in the bound d on j.

Provided F is analytic, this property is inherited by the functions

$$y_j = y_j(x_0, x_1, \ldots, x_j) \in I\!R^m,$$

and their derivatives satisfy the identities

$$\frac{\partial y_j}{\partial x_i} = \frac{\partial y_{j-i}}{\partial x_0} = A_{j-i}(x_0, x_1, \ldots, x_{j-i}) \tag{8}$$

as established in [7]. This yields in particular

$$\begin{aligned}
\frac{\partial y_0}{\partial x_0} = \frac{\partial y_1}{\partial x_1} = \frac{\partial y_2}{\partial x_2} = \frac{\partial y_3}{\partial x_3} &= A_0 = F'(x_0) \\
\frac{\partial y_1}{\partial x_0} = \frac{\partial y_2}{\partial x_1} = \frac{\partial y_3}{\partial x_2} &= A_1 = F''(x_0)x_1 \\
\frac{\partial y_2}{\partial x_0} = \frac{\partial y_3}{\partial x_1} &= A_2 = F''(x_0)x_2 + \frac{1}{2}F'''(x_0)x_1x_1 \\
\frac{\partial y_3}{\partial x_0} &= A_3 = F''(x_0)x_3 + F'''(x_0)x_1x_2 + \frac{1}{6}F^{(4)}(x_0)x_1x_1x_1 \\
&\dots
\end{aligned}$$

The $m \times n$ matrices $A_k, k = 0, \dots, d$, are actually the Taylor coefficients of the Jacobian path $F'(x(t))$, a fact that is of interest primarily in the context of ordinary differential equations and differential algebraic equations.

Given the tape of an active section and the coefficients x_j, the resulting y_j and their derivatives A_j can be evaluated by appropriate calls to the ADOL-C forward mode implementations and the ADOL-C reverse mode implementations. The scalar versions of the forward mode propagate just one truncated Taylor series from the $(x_j)_{j \le d}$ to the $(y_j)_{j \le d}$. The vector versions of the forward mode propagate families of $p \ge 1$ such truncated Taylor series in order to reduce the relative cost of the overhead incurred in the tape interpretation. In detail, ADOL-C provides

```
int zos_forward(tag,m,n,keep,x,y)
short int tag;          // tape identification
int m;                  // number of dependent variables m
int n;                  // number of independent variables n
int keep;               // flag for reverse mode preparation
double x[n];            // independent vector x = x_0
double y[m];            // dependent vector y = F(x_0)
```

for the **z**ero-**o**rder **s**calar forward mode. This driver computes $y = F(x)$ with $0 \le \mathsf{keep} \le 1$. The integer flag keep plays a similar role as in the call to trace_on: It determines if zos_forward writes the first Taylor coefficients of all intermediate quantities into a buffered temporary file, i.e., the value stack, in preparation for a subsequent reverse mode evaluation. The value $\mathsf{keep} = 1$ prepares for fos_reverse or fov_reverse as exlained below.

To compute first-order derivatives, one has

```
int fos_forward(tag,m,n,keep,x0,x1,y0,y1)
short int tag;          // tape identification
int m;                  // number of dependent variables m
int n;                  // number of independent variables n
```

```
int keep;                    // flag for reverse mode preparation
double x0[n];                // independent vector x_0
double x1[n];                // tangent vector x_1
double y0[m];                // dependent vector y_0 = F(x_0)
double y1[m];                // first derivative y_1 = F'(x_0)x_1
```

for the **f**irst-**o**rder **s**calar forward mode. Here, one has $0 \leq \mathsf{keep} \leq 2$, where

$$\mathsf{keep} = \begin{cases} 1 & \text{prepares for } \mathsf{fos_reverse} \text{ or } \mathsf{fov_reverse} \\ 2 & \text{prepares for } \mathsf{hos_reverse} \text{ or } \mathsf{hov_reverse} \end{cases}$$

as exlained below. For the **f**irst-**o**rder **v**ector forward mode, ADOL-C provides

```
int fov_forward(tag,m,n,p,x0,X,y0,Y)
short int tag;               // tape identification
int m;                       // number of dependent variables m
int n;                       // number of independent variables n
int p;                       // number of directions
double x0[n];                // independent vector x_0
double X[n][p];              // tangent matrix X
double y0[m];                // dependent vector y_0 = F(x_0)
double Y[m][p];              // first derivative matrix Y = F'(x)X
```

For the computation of higher derivative, the driver

```
int hos_forward(tag,m,n,d,keep,x0,X,y0,Y)
short int tag;               // tape identification
int m;                       // number of dependent variables m
int n;                       // number of independent variables n
int d;                       // highest derivative degree d
int keep;                    // flag for reverse mode preparation
double x0[n];                // independent vector x_0
double X[n][d];              // tangent matrix X
double y0[m];                // dependent vector y_0 = F(x_0)
double Y[m][d];              // derivative matrix Y
```

implementing the **h**igher-**o**rder **s**calar forward mode. The rows of the matrix X must correspond to the independent variables in the order of their initialization by the $\ll=$ operator. The columns of $X = \{x_j\}_{j=1\ldots d}$ represent Taylor coefficient vectors as in (4). The rows of the matrix Y must correspond to the dependent variables in the order of their selection by the $\gg=$ operator. The columns of $Y = \{y_j\}_{j=1\ldots d}$ represent Taylor coefficient vectors as in (6), i.e., hos_forward computes the values $y_0 = F(x_0)$, $y_1 = F'(x_0)x_1$, ...,

where $X = [x_1, x_2, \ldots, x_d]$ and $Y = [y_1, y_2, \ldots, y_d]$. Furthermore, one has $0 \leq$ keep $\leq d+1$, with

$$\text{keep} \begin{cases} = 1 & \text{prepares for fos_reverse or fov_reverse} \\ > 1 & \text{prepares for hos_reverse or hov_reverse} \end{cases}$$

Once more, there is also a vector version given by

```
int hov_forward(tag,m,n,d,p,x0,X,y0,Y)
short int tag;                  // tape identification
int m;                          // number of dependent variables m
int n;                          // number of independent variables n
int d;                          // highest derivative degree d
int p;                          // number of directions p
double x0[n];                   // independent vector x_0
double X[n][p][d];              // tangent matrix X
double y0[m];                   // dependent vector y_0 = F(x_0)
double Y[m][p][d];              // derivative matrix Y
```

for the **h**igher-**o**rder **v**ector forward mode that computes $y_0 = F(x_0)$, $Y_1 = F'(x_0)X_1$, ..., where $X = [X_1, X_2, \ldots, X_d]$ and $Y = [Y_1, Y_2, \ldots, Y_d]$.

There are also overloaded versions providing a general forward-call. Details of the appropriate calling sequences are given in Section 6.

Once, the required information is generated due to a forward mode evaluation with an approriate value of the parameter keep, one may use the following implementation variants of the reverse mode. To compute first-order derivatives one can use

```
int fos_reverse(tag,m,n,u,z)
short int tag;                  // tape identification
int m;                          // number of dependent variables m
int n;                          // number of independent variables n
double u[m];                    // weight vector u
double z[m];                    // resulting adjoint value z^T = u^T F'(x)
```

as **f**irst-**o**rder **s**calar reverse mode implementation that computes the product $z^T = u^T F'(x)$ after calling zos_forward, fos_forward, or hos_forward with keep=1. The corresponding **f**irst-**o**rder **v**ector reverse mode driver is given by

```
int fov_reverse(tag,m,n,q,U,Z)
short int tag;                  // tape identification
int m;                          // number of dependent variables m
int n;                          // number of independent variables n
```

```
int q;                          // number of weight vectors q
double U[q][m];                 // weight matrix U
double Z[q][m];                 // resulting adjoint Z = UF'(x)
```

that can be used after calling zos_forward, fos_forward, or hos_forward with keep=1. To compute higher-order derivatives, ADOL-C provides

```
int hos_reverse(tag,m,n,d,u,Z)
short int tag;                  // tape identification
int m;                          // number of dependent variables m
int n;                          // number of independent variables n
int d;                          // highest derivative degree d
double u[m];                    // weight vector u
double Z[m][d+1];               // resulting adjoints
```

as **h**igher-**o**rder **s**calar reverse mode implementation yielding the adjoints $z_0^T = u^T F'(x_0) = u^T A_0$, $z_1^T = u^T F''(x_0)x_1 = u^T A_1$, ..., where $Z = [z_0, z_1, \ldots, z_d]$ after calling fos_forward or hos_forward with keep $= d + 1 > 1$. The vector version is given by

```
int hov_reverse(tag,m,n,d,q,U,Z,nz)
short int tag;                  // tape identification
int m;                          // number of dependent variables m
int n;                          // number of independent variables n
int d;                          // highest derivative degree d
double U[q][m];                 // weight vector u
double Z[q][m][d+1];            // resulting adjoints
short int nz[q][n];             // nonzero pattern of Z
```

as **h**igher-**o**rder **v**ector reverse mode driver to compute the adjoints $Z_0 = UF'(x_0) = UA_0$, $Z_1 = UF''(x_0)x_1 = UA_1$, ..., where $Z = [Z_0, Z_1, \ldots, Z_d]$ after calling fos_forward or hos_forward with keep $= d + 1 > 1$. After the function call, the last argument of hov_reverse contains information about the sparsity pattern, i.e. each nz[i][j] has a value that characterizes the functional relation between the i-th component of $UF'(x)$ and the j-th independent value x_j as:

0	trivial	2	polynomial	4	transcendental
1	linear	3	rational	5	non-smooth

Here, "trivial" means that there is no dependence at all and "linear" means that the partial derivative is a constant that does not dependent on other variables either. "Non-smooth" means that one of the functions on the path between x_i and y_j was evaluated at a point where it is not differentiable. All positive labels $1, 2, 3, 4, 5$ are pessimistic in that the actual functional relation may in fact be simpler, for example due to exact cancellations.

There are also overloaded versions providing a general reverse-call. Details of the appropriate calling sequences are given in the following Section 6.

6 Overloaded Forward and Reverse Calls

In this section, the several versions of the forward and reverse routines, which utilize the overloading capabilities of C++, are described in detail. With exception of the bit pattern versions all interfaces are prototyped in the header file `<adolc/interfaces.h>`, where also some more specialized forward and reverse routines are explained. Furthermore, ADOL-C provides C and Fortran-callable versions prototyped in the same header file. The bit pattern versions of forward and reverse introduced in the subsection 6.3 are prototyped in the header file `<adolc/sparse/sparse.h>`, which will be included by the header file `<adolc/interfaces.h>` automatically.

6.1 The Scalar Case

Given any correct tape, one may call from within the generating program, or subsequently during another run, the following procedure:

```
int forward(tag,m,n,d,keep,X,Y)
short int tag;              // tape identification
int m;                      // number of dependent variables m
int n;                      // number of independent variables n
int d;                      // highest derivative degree d
int keep;                   // flag for reverse sweep
double X[n][d+1];           // Taylor coefficients X of independent variables
double Y[m][d+1];           // Taylor coefficients Y as in (6)
```

The rows of the matrix X must correspond to the independent variables in the order of their initialization by the $\ll=$ operator. The columns of $X = \{x_j\}_{j=0\ldots d}$ represent Taylor coefficient vectors as in (4). The rows of the matrix Y must correspond to the dependent variables in the order of their selection by the $\gg=$ operator. The columns of $Y = \{y_j\}_{j=0\ldots d}$ represent Taylor coefficient vectors as in (6). Thus the first column of Y contains the function value $F(x)$ itself, the next column represents the first Taylor coefficient vector of F, and the last column the d-th Taylor coefficient vector. The integer flag keep determines how many Taylor coefficients of all intermediate quantities are written into the value stack as explained in 5. If keep is omitted, it defaults to 0.

The given tag value is used by forward to determine the name of the file on which the tape was written. If the tape file does not exist, forward assumes that the relevant tape is still in core and reads from the buffers. After the execution of an active section with

keep $= 1$ or a call to forward with any keep $\leq d+1$, one may call the function reverse with d $=$ keep $- 1$ and the same tape identifier tag. When u is a vector and Z an $n \times (d+1)$ matrix reverse is executed in the scalar mode by the calling sequence

```
int reverse(tag,m,n,d,u,Z)
short int tag;          // tape identification
int m;                  // number of dependent variables m
int n;                  // number of independent variables n
int d;                  // highest derivative degree d
double u[m];            // weighting vector u
double Z[n][d+1];       // resulting adjoints Z
```

to compute the adjoints $z_0^T = u^T F'(x_0) = u^T A_0$, $z_1^T = u^T F''(x_0)x_1 = u^T A_1$, ..., where $Z = [z_0, z_1, \ldots, z_d]$.

6.2 The Vector Case

When U is a matrix reverse is executed in the vector mode by the following calling sequence

```
int reverse(tag,m,n,d,q,U,Z,nz)
short int tag;          // tape identification
int m;                  // number of dependent variables m
int n;                  // number of independent variables n
int d;                  // highest derivative degree d
int q;                  // number of weight vectors q
double U[q][m];         // weight matrix U
double Z[q][n][d+1];    // resulting adjoints
short nz[q][n];         // nonzero pattern of Z
```

to compute the adjoints $Z_0 = UF'(x_0) = UA_0$, $Z_1 = UF''(x_0)x_1 = UA_1$, ..., where $Z = [Z_0, Z_1, \ldots, Z_d]$. When the arguments p and U are omitted, they default to m and the identity matrix of order m, respectively.

Through the optional argument nz of reverse one can compute information about the sparsity pattern of Z as described in detail in the previous Section 5.

The return values of reverse calls can be interpreted according to Table 1, but negative return values are not valid, since the corresponding forward sweep would have stopped without completing the necessary taylor file. The return value of reverse may be higher than that of the preceding forward call because some operations that were evaluated at a critical argument during the forward sweep were found not to impact the dependents during the reverse sweep.

In both scalar and vector mode, the degree d must agree with keep $- 1$ for the most recent call to forward, or it must be equal to zero if reverse directly follows the taping of an active section. Otherwise, reverse will return control with a suitable error message. In order to avoid possible confusion, the first four arguments must always be present in the calling sequence. However, if m or d attain their trivial values 1 and 0, respectively, then corresponding dimensions of the arrays X, Y, u, U, or Z can be omitted, thus eliminating one level of indirection. For example, we may call reverse(tag,1,n,0,1.0,g) after declaring double g[n] to calculate a gradient of a scalar-valued function.

Sometimes it may be useful to perform a forward sweep for families of Taylor series with the same leading term. This vector version of forward can be called in the form

```
int forward(tag,m,n,d,p,x0,X,y0,Y)
short int tag;              // tape identification
int m;                      // number of dependent variables m
int n;                      // number of independent variables n
int d;                      // highest derivative degree d
int p;                      // number of Taylor series p
double x0[n];               // values of independent variables x_0
double X[n][p][d];          // Taylor coefficients X of independent variables
double y0[m];               // values of dependent variables y_0
double Y[m][p][d];          // Taylor coefficients Y of dependent variables
```

where X and Y hold the Taylor coefficients of first and higher degree and x0, y0 the common Taylor coefficients of degree 0. There is no option to keep the values of active variables that are going out of scope or that are overwritten. Therefore this function cannot prepare a subsequent reverse sweep. The return integer serves as a flag to indicate quadratures or altered comparisons as described above in Section 2.6.

Since the calculation of Jacobians is probably the most important automatic differentiation task, we have provided a specialization of vector forward to the case where $d = 1$. This version can be called in the form

```
int forward(tag,m,n,p,x,X,y,Y)
short int tag;              // tape identification
int m;                      // number of dependent variables m
int n;                      // number of independent variables n
int p;                      // number of partial derivatives p
double x[n];                // values of independent variables x_0
double X[n][p];             // seed derivatives of independent variables X
double y[m];                // values of dependent variables y_0
double Y[m][p];             // first derivatives of dependent variables Y
```

When this routine is called with p = n and X the identity matrix, the resulting Y is simply the Jacobian $F'(x_0)$. In general, one obtains the $m \times p$ matrix $Y = F'(x_0)\,X$ for the chosen

initialization of X. In a workstation environment a value of p somewhere between 10 and 50 appears to be fairly optimal. For smaller p the interpretive overhead is not appropriately amortized, and for larger p the p-fold increase in storage causes too many page faults. Therefore, large Jacobians that cannot be compressed via column coloring as could be done for example using the driver sparse_jac should be "strip-mined" in the sense that the above first-order-vector version of forward is called repeatedly with the successive $n \times p$ matrices X forming a partition of the identity matrix of order n.

6.3 Dependence Analysis

The sparsity pattern of Jacobians is often needed to set up data structures for their storage and factorization or to allow their economical evaluation by compression [3]. Compared to the evaluation of the full Jacobian $F'(x_0)$ in real arithmetic computing the Boolean matrix $\tilde{P} \in \{0,1\}^{m \times n}$ representing its sparsity pattern in the obvious way requires a little less run-time and certainly a lot less memory.

The entry $\tilde{P}_{ji}$ in the j-th row and i-th column of $\tilde{P}$ should be $1 = true$ exactly when there is a data dependence between the i-th independent variable x_i and the j-th dependent variable y_j. Just like for real arguments one would wish to compute matrix-vector and vector-matrix products of the form $\tilde{P}\tilde{v}$ or $\tilde{u}^T\tilde{P}$ by appropriate forward and reverse routines where $\tilde{v} \in \{0,1\}^n$ and $\tilde{u} \in \{0,1\}^m$. Here, multiplication corresponds to logical AND and addition to logical OR, so that algebra is performed in a semi-ring.

For practical reasons it is assumed that $s = 8*$sizeof(unsigned long int) such Boolean vectors $\tilde{v}$ and $\tilde{u}$ are combined to integer vectors $v \in \mathbb{N}^n$ and $u \in \mathbb{N}^m$ whose components can be interpreted as bit patterns. Moreover p or q such integer vectors may be combined column-wise or row-wise to integer matrices $X \in \mathbb{N}^{n \times p}$ and $U \in \mathbb{N}^{q \times m}$, which naturally correspond to Boolean matrices $\tilde{X} \in \{0,1\}^{n \times (sp)}$ and $\tilde{U} \in \{0,1\}^{(sq) \times m}$. The provided bit pattern versions of forward and reverse allow to compute integer matrices $Y \in \mathbb{N}^{m \times p}$ and $Z \in \mathbb{N}^{q \times m}$ corresponding to

$$\tilde{Y} = \tilde{P}\tilde{X} \qquad \text{and} \qquad \tilde{Z} = \tilde{U}\tilde{P}\,, \tag{9}$$

respectively, with $\tilde{Y} \in \{0,1\}^{m \times (sp)}$ and $\tilde{U} \in \{0,1\}^{(sq) \times n}$. In general, the application of the bit pattern versions of forward or reverse can be interpreted as propagating dependences between variables forward or backward, therefore both the propagated integer matrices and the corresponding Boolean matrices are called *dependence structures*.

The bit pattern forward routine

```
int forward(tag,m,n,p,x,X,y,Y,mode)
short int tag;          // tape identification
int m;                  // number of dependent variables m
int n;                  // number of independent variables n
```

```
int p;                          // number of integers propagated p
double x[n];                    // values of independent variables x_0
unsigned long int X[n][p];      // dependence structure X
double y[m];                    // values of dependent variables y_0
unsigned long int Y[m][p];      // dependence structure Y according to (9)
char mode;                      // 0 : safe mode (default), 1 : tight mode
```

can be used to obtain the dependence structure Y for a given dependence structure X. The dependence structures are represented as arrays of unsigned long int the entries of which are interpreted as bit patterns as described above. For example, for $n = 3$ the identity matrix I_3 should be passed with $p = 1$ as the 3×1 array

$$\mathsf{X} \;=\; \begin{pmatrix} 10000000\ 00000000\ 00000000\ 00000000_2 \\ 01000000\ 00000000\ 00000000\ 00000000_2 \\ 00100000\ 00000000\ 00000000\ 00000000_2 \end{pmatrix}$$

in the 4-byte long integer format. The parameter mode determines the mode of dependence analysis as explained already in subsection 4.3.

A call to the corresponding bit pattern reverse routine

```
int reverse(tag,m,n,q,U,Z,mode)
short int tag;                  // tape identification
int m;                          // number of dependent variables m
int n;                          // number of independent variables n
int q;                          // number of integers propagated q
unsigned long int U[q][m];      // dependence structure U
unsigned long int Z[q][n];      // dependence structure Z according to (9)
char mode;                      // 0 : safe mode (default), 1 : tight mode
```

yields the dependence structure Z for a given dependence structure U.

To determine the whole sparsity pattern $\tilde{P}$ of the Jacobian $F'(x)$ as an integer matrix P one may call forward or reverse with $p \geq n/s$ or $q \geq m/s$, respectively. For this purpose the corresponding dependence structure X or U must be defined to represent the identity matrix of the respective dimension. Due to the fact that always a multiple of s Boolean vectors are propagated there may be superfluous vectors, which can be set to zero.

The return values of the bit pattern forward and reverse routines correspond to those described in Table 1.

One can control the storage growth by the factor p using "strip-mining" for the calls of forward or reverse with successive groups of columns or respectively rows at a time, i.e. partitioning X or U appropriately as described for the computation of Jacobians in Section 6.2.

7 Tapeless forward differentiation in ADOL-C

Up to version 1.9.0, the development of the ADOL-C software package was based on the decision to store all data necessary for derivative computation on tapes, where large applications require the tapes to be written out to corresponding files. In almost all cases this means a considerable drawback in terms of run time due to the excessive memory accesses. Using these tapes enables ADOL-C to offer multiple functions. However, it is not necessary for all tasks of derivative computation to do that.

Starting with version 1.10.0, ADOL-C now features a tapeless forward mode for computing first order derivatives in scalar mode, i.e., $\dot{y} = F'(x)\dot{x}$, and in vector mode, i.e., $\dot{Y} = F'(x)\dot{X}$. This tapeless variant coexists with the more universal tape based mode in the package. The following subsections describe the source code modifications required to use the tapeless forward mode of ADOL-C.

7.1 Modifying the Source Code

Let us consider the coordinate transformation from Cartesian to spherical polar coordinates given by the function $F : \mathbb{R}^3 \to \mathbb{R}^3$, $y = F(x)$, with

$$y_1 = \sqrt{x_1^2 + x_2^2 + x_3^2}, \qquad y_2 = \arctan\left(\sqrt{x_1^2 + x_2^2}/x_3\right), \qquad y_3 = \arctan(x_2/x_1),$$

as an example. The corresponding source code is shown in Figure 5. Changes to the source code that are necessary for applying the tapeless forward ADOL-C are described in the following two subsections, where the vector mode version is described as extension of the scalar mode.

The scalar mode

To use the tapeless forward mode, one has to include one of the header files `adolc/adolc.h` or `adolc/adouble.h` where the latter should be preferred since it does not include the tape based functions defined in other header files. Hence, including `adolc/adouble.h` avoids mode mixtures, since `adolc/adolc.h` is just a wrapper for including all public headers of the ADOL-C package and does not offer own functions. Since the two ADOL-C forward mode variants tape-based and tapeless, are prototyped in the same header file, the compiler needs to know if a tapeless version is intended. This can be done by defining a preprocessor macro named ADOLC_TAPELESS. Note that it is important to define this macro before the header file is included. Otherwise, the tape-based version of ADOL-C will be used.

As in the tape based forward version of ADOL-C all derivative calculations are introduced by calls to overloaded operators. Therefore, similar to the tape-based version all independent, intermediate and dependent variables must be declared with type adouble.

```
#include <iostream>
using namespace std;

int main() {
   double x[3], y[3];

   for (int i=0; i<3; ++i)                    // Initialize x_i
      ...

   y[0] = sqrt(x[0]*x[0]+x[1]*x[1]+x[2]*x[2]);
   y[1] = atan(sqrt(x[0]*x[0]+x[1]*x[1])/x[2]);
   y[2] = atan(x[1]/x[0]);

   cout << "y1=" << y[0] << " , y2=" << y[1] << " , y3=" << y[2] << endl;

   return 0;
}
```

Figure 5: Example for tapeless forward mode

The whole tapeless functionality provided by `adolc/adolc.h` was written as complete inline intended code due to run time aspects, where the real portion of inlined code can be influenced by switches for many compilers. Likely, the whole derivative code is inlined by default. Our experiments with the tapeless mode have produced complete inlined code by using standard switches (optimization) for GNU and Intel C++ compiler.

To avoid name conflicts resulting from the inlining the tapeless version has its own namespace `adtl`. As a result four possibilities of using the adouble type are available for the tapeless version:

- Defining a new type

  ```
  typedef adtl::adouble adouble;
  ...
  adouble tmp;
  ```

  This is the preferred way. Remember, you can not write an own adouble type/class with different meaning after doing the typedef.

- Declaring with namespace prefix

  ```
  adtl::adouble tmp;
  ```

Not the most handsome and efficient way with respect to coding but without any doubt one of the safest ways. The identifier adouble is still available for user types/classes.

- Trusting macros

```
#define adouble adtl::adouble
...
adouble tmp;
```

This approach should be used with care, since standard defines are text replacements.

- Using the complete namespace

```
#using namespace adtl;
...
adouble tmp;
```

A very clear approach with the disadvantage of uncovering all the hidden secrets. Name conflicts may arise!

After defining the variables only two things are left to do. First one needs to initialize the values of the independent variables for the function evaluation. This can be done by assigning the variables a double value. The ad-value is set to zero in this case. Additionally, the tapeless forward mode variant of ADOL-C offers a function named setValue for setting the value without changing the ad-value. To set the ad-values of the independent variables ADOL-C offers two possibilities:

- Using the constructor

```
adouble x1(2,1), x2(4,0), y;
```

This would create three adoubles x_1, x_2 and y. Obviously, the latter remains uninitialized. In terms of function evaluation x_1 holds the value 2 and x_2 the value 4 whereas the derivative values are initialized to $\dot{x}_1 = 1$ and $\dot{x}_2 = 0$.

- Setting point values directly

```
adouble x1=2, x2=4, y;
...
x1.setADValue(1);
x2.setADValue(0);
```

The same example as above but now using setADValue-method for initializing the derivative values.

The derivatives can be obtained at any time during the evaluation process by calling the getADValue-method

```
adouble y;
...
cout << y.getADValue();
```

Fig. 6 shows the resulting source code incorporating all required changes for the example given above.

```
#include <iostream>
using namespace std;

#define ADOLC_TAPELESS
#include <adolc/adouble.h>
typedef adtl::adouble adouble;

int main() {
   adouble x[3], y[3];

   for (int i=0; i<3; ++i)                    // Initialize x_i
      ...

   x[0].setADValue(1);                       // derivative of f with respect to x_1

   y[0] = sqrt(x[0]*x[0]+x[1]*x[1]+x[2]*x[2]);
   y[1] = atan(sqrt(x[0]*x[0]+x[1]*x[1])/x[2]);
   y[2] = atan(x[1]/x[0]);

   cout << "y1=" << y[0].getValue() << " , y2=" << y[1].getValue ... ;
   cout << "dy2/dx1 = " << y[1].getADValue() << endl;

   return 0;
}
```

Figure 6: Example for tapeless scalar forward mode

The vector mode

In scalar mode only one direction element has to be stored per adouble whereas a field of p elements is needed in the vector mode to cover the computations for the given p directions.

The resulting changes to the source code are described in this section.

Similar to tapeless scalar forward mode, the tapeless vector forward mode is used by defining ADOLC_TAPELESS. Furthermore, one has to define an additional preprocessor macro named NUMBER_DIRECTIONS. This macro takes the maximal number of directions to be used within the resulting vector mode. Just as ADOLC_TAPELESS the new macro must be defined before including the `<adolc.h/adouble.h>` header file since it is ignored otherwise.

In many situations recompiling the source code to get a new number of directions is at least undesirable. ADOL-C offers a function named setNumDir to work around this problem partially. Calling this function, ADOL-C does not take the number of directions from the macro NUMBER_DIRECTIONS but from the argument of setNumDir. A corresponding source code would contain the following lines:

```
#define NUMBER_DIRECTIONS 10
...
adtl::setNumDir(5);
```

Note that using this function does not change memory requirements that can be roughly determined by (NUMBER_DIRECTIONS+1)*(number of adoubles).

Compared to the scalar case setting and getting the derivative values, i.e. the directions, is more involved. Instead of working with single double values, pointer to fields of doubles are used as illustrated by the following example:

```
#define NUMBER_DIRECTIONS 10
...
adouble x, y;
double *ptr=new double[NUMBER_DIRECTIONS];
...
x1=2;
x1.setADValue(ptr);
...
ptr=y.getADValue();
```

Additionally, the tapeless vector forward mode of ADOL-C offers two new methods for setting/getting the derivative values. Similar to the scalar case, double values are used but due to the vector mode the position of the desired vector element must be supplied in the argument list:

```
#define NUMBER_DIRECTIONS 10
...
adouble x, y;
...
x1=2;
x1.setADValue(5,1);                    // set the 6th point value of x to 1.0
...
cout << y.getADValue(3) << endl;       // print the 4th derivative value of y
```

The resulting source code containing all changes that are required is shown in Fig. 7

7.2 Compiling and Linking the Source Code

After incorporating the required changes, one has to compile the source code and link the object files to get the executable. As long as the ADOL-C header files are not included in the absolute path the compile sequence should be similar to the following example:

```
g++ -I/home/username/adolc_base/include -c tapeless_scalar.cpp
```

The `-I` option tells the compiler where to search for the ADOL-C header files. This option can be omitted when the headers are included with absolute path or if ADOL-C is installed in a "global" directory.

Since the tapeless forward version of ADOL-C is implemented in the header `adouble.h` as complete inline intended version, the object files do not need to be linked against any external ADOL-C code or the ADOL-C library. Therefore, the example started above could be finished with the following command:

```
g++ -o tapeless_scalar tapeless_scalar.o
```

The mentioned source codes tapeless_scalar.c and tapeless_vector.c illustrating the use of the for tapeless scalar and vector mode can be found in the directory examples.

7.3 Concluding Remarks for the Tapeless Forward Mode Variant

As many other AD methods the tapeless forward mode provided by the ADOL-C package has its own strengths and drawbacks. Please read the following section carefully to become familiar with the things that can occur:

- Advantages:
  - Code speed
    Increasing computation speed was one of the main aspects in writing the tapeless code. In many cases higher performance can be expected this way.

```
#include <iostream>
using namespace std;

#define ADOLC_TAPELESS
#define NUMBER_DIRECTIONS 3
#include <adolc/adouble.h>
typedef adtl::adouble adouble;

int main() {
   adouble x[3], y[3];

   for (int i=0; i<3; ++i) {
      ...                    // Initialize x_i
      for (int j=0; j<3; ++j) if (i==j) x[i].setADValue(j,1);
   }

   y[0] = sqrt(x[0]*x[0]+x[1]*x[1]+x[2]*x[2]);
   y[1] = atan(sqrt(x[0]*x[0]+x[1]*x[1])/x[2]);
   y[2] = atan(x[1]/x[0]);

   cout << "y1=" << y[0].getValue() << " , y2=" << y[1].getValue ... ;
   cout << "jacobian : " << endl;
   for (int i=0; i<3; ++i) {
      for (int j=0; j<3; ++j)
         cout << y[i].getADValue(j) << " ";
      cout << endl;
   }

   return 0;
}
```

Figure 7: Example for tapeless vector forward mode

- Easier linking process
  As another result from the code inlining the object code does not need to be linked against an ADOL-C library.
- Smaller overall memory requirements
  Tapeless ADOL-C does not write tapes anymore, as the name implies. Loop "unrolling" can be avoided this way. Considered main memory plus disk space as overall memory requirements the tapeless version can be executed in a more

efficient way.

- Drawbacks:
  - Main memory limitations
    The ability to compute derivatives to a given function is bounded by the main memory plus swap size when using tapeless ADOL-C. Computation from swap should be avoided anyway as far as possible since it slows down the computing time drastically. Therefore, if the program execution is terminated without error message insufficient memory size can be the reason among other things. The memory requirements M can be determined roughly as followed:
    * Scalar mode: $M =$(number of adoubles)$*2 + M_p$
    * Vector mode: $M =$(number of adoubles)*(NUMBER_DIRECTIONS+1) $+ M_p$

    where the storage size of all non adouble based variables is described by M_p.
  - Compile time
    As discussed in the previous sections, the tapeless forward mode of the ADOL-C package is implemented as inline intended version. Using this approach results in a higher source code size, since every operation involving at least one adouble stands for the operation itself as well as for the corresponding derivative code after the inlining process. Therefore, the compilation time needed for the tapeless version may be higher than that of the tape based code.
  - Code Size
    A second drawback and result of the code inlining is the increase of code sizes for the binaries. The increase factor compared to the corresponding tape based program is difficult to quantify as it is task dependent. Practical results have shown that values between 1.0 and 2.0 can be expected. Factors higher than 2.0 are possible too and even values below 1.0 have been observed.

8 Installing and Using ADOL-C

8.1 Generating the ADOL-C Library

The currently built system is best summarized by the ubiquitous gnu install triplet

`configure - make - make install` .

Executing this three steps from the package base directory `</SOMEPATH/adolc-1.9.0>` will compile the static and the dynamic ADOL-C library with default options and install the package (libraries and headers) into the default installation directory `<$HOME/adolc_base>`.

Before doing so the user may modify the header file `usrparms.h` in order to tailor the ADOL-C package to his needs in the particular system environment as discussed in Section

3.2. The configure procedure which creates the necessary `Makefile`s can be customized by use of some switches. Available options and their meaning can be obtained by executing `./configure --help` from the package base directory.

All object files and other intermediately generated files can be removed by the call `make clean`. Uninstalling ADOL-C by executing `make uninstall` is only reasonable after a previous called `make install` and will remove all installed package files but will leave the created directories behind.

8.2 Compiling and Linking the Example Programs

The installation procedure described in Subsection 8.1 also provides the `Makefile`s necessary to compile the example programs placed in the directories `<adolc-1.9.0>/examples` and `<adolc-1.9.0>/examples/additional_examples`. However, one must `configure` the ADOL-C package to enable the compilation of examples by supplying appropriate options. Available options are:

`--with-docexa`	build all examples discussed in this manual (compare Section 9)
`--with-addexa`	build all additional examples (See file `README` in the various subdirectories)

Just calling `make` from the packages base directory generates all configured examples and the library if necessary. Compiling from subdirectory `examples` or one of its subfolders is possible too. At least one kind of the ADOL-C library (static or shared) must have been built previously in that case. Hence, building the library is always the first step.

8.3 Description of Important Header Files

The application of the facilities of ADOL-C requires the user source code (program or module) to include appropriate header files where the desired data types and routines are prototyped. The new hierarchy of header files enables the user to take one of two possible ways to access the right interfaces. The first and easy way is recommended to beginners: As indicated in Table 3 the provided *global* header file `<adolc/adolc.h>` can be included by any user code to support all capabilities of ADOL-C depending on the particular programming language of the source.

The second way is meant for the more advanced ADOL-C user: Some source code includes only those interfaces used by the particular application. The respectively needed header files are indicated throughout the manual. Existing application determined dependences between the provided ADOL-C routines are realized by automatic includes of headers in order to maintain easy use. The header files important to the user are described in the Tables 4 and 5.

`<adolc/adolc.h>`	→ global header file available for easy use of ADOL-C; • includes all ADOL-C header files depending on whether the users source is C++ or C code.
`<adolc/usrparms.h>`	→ user customization of ADOL-C package (see Section 3.2); • after a change of user options the ADOL-C library `libadolc.*` has to be rebuilt (see Section 8.1); • is included by all ADOL-C header files and thus by all user programs.

Table 3: Global header files

`<adolc/adouble.h>`	→ provides the interface to the basic active scalar data type of ADOL-C: class adouble (see Section 2);
`<adolc/taputil.h>`	→ provides functions to start/stop the tracing of active sections (see Sections 2.2) as well as utilities to obtain tape statistics (see Sections 3.1); • is included by the header `<adolc/adouble.h>`.

Table 4: Important header files: tracing/taping

8.4 Compiling and Linking C/C++ Programs

To compile a C/C++ program or single module using ADOL-C data types and routines one has to ensure that all necessary header files according to Section 8.3 are included. All modules involving *active* data types as adouble have to be compiled as C++. Modules that make use of a previously generated tape to evaluate derivatives can either be programmed in ANSI-C (while avoiding all C++ interfaces) or in C++. Depending on the chosen programming language the header files provide the right ADOL-C prototypes. For linking the resulting object codes the library `libadolc.*` must be used (see Section 8.1).

8.5 Adding Quadratures as Special Functions

Suppose an integral

$$f(x) = \int_0^x g(t)dt$$

is evaluated numerically by a user-supplied function

```
double myintegral(double& x);
```

Similarly, let us suppose that the integrand itself is evaluated by a user-supplied block of C code integrand, which computes a variable with the fixed name val from a variable with the fixed name arg. In many cases of interest, integrand will simply be of the form

`<adolc/interfaces.h>`	→ provides interfaces to the forward and reverse routines as basic versions of derivative evaluation (see Section 6); • comprises C++, C, and Fortran-callable versions; • includes the header `<adolc/sparse/sparse.h>`; • is included by the header `<adolc/drivers/odedrivers.h>`.
`<adolc/drivers.h>`	→ provides "easy to use" drivers for solving optimization problems and nonlinear equations (see Section 4.1); • comprises C and Fortran-callable versions.
`<adolc/sparse/ sparse.h>`	→ provides the "easy to use" sparse drivers to exploit the sparsity structure of Jacobians (see Section 4.3); → provides interfaces to C++-callable versions of forward and reverse routines propagating bit patterns (see Section 6.3); • is included by the header `<adolc/interfaces.h>`.
`<adolc/sparse/ jacutils.h>`	→ provides interfaces to the underlying C-callable versions of forward and reverse routines propagating bit patterns.
`<adolc/drivers/ odedrivers.h>`	→ provides "easy to use" drivers for numerical solution of ordinary differential equations (see Section 4.2); • comprises C++, C, and Fortran-callable versions; • includes the header `<adolc/interfaces.h>`.
`<adolc/drivers/ taylor.h>`	→ provides "easy to use" drivers for evaluation of higher order derivative tensors (see Section 4.4) and inverse/implicit function differentiation (see Section 4.5); • comprises C++ and C-callable versions.
`<adolc/adalloc.h>`	→ provides C++ and C functions for allocation of vectors, matrices and three dimensional arrays of doubles.

Table 5: Important header files: evaluation of derivatives

$$\{\ \mathsf{val} = \mathsf{expression}(\mathsf{arg})\ \}\ .$$

In general, the final assignment to val may be preceded by several intermediate calculations, possibly involving local active variables of type adouble, but no external or static variables of that type. However, integrand may involve local or global variables of type double or int, provided they do not depend on the value of arg. The variables arg and val are declared automatically; and as integrand is a block rather than a function, integrand should have no header line.

Now the function myintegral can be overloaded for adouble arguments and thus included in the library of elementary functions by the following modifications:

1. At the end of the file `<adolc/adouble.cpp>`, include the full code defining double myintegral(double& x), and add the line

   extend_quad(myintegral, integrand);

   This macro is extended to the definition of adouble myintegral(adouble& arg). Then remake the library `libadolc.*` (see Section 8.1).

2. In the definition of the class ADOLC_DLL_EXPORT adouble in <adolc/adouble.h>, add the statement

friend adouble myintegral(adouble&).

In the first modification, myintegral represents the name of the double function, whereas integrand represents the actual block of C code.

For example, in case of the inverse hyperbolic cosine, we have myintegral = acosh. Then integrand can be written as { val = sqrt(arg*arg-1); } so that the line

extend_quad(acosh,val = sqrt(arg*arg-1));

can be added to the file <adolc/adouble.cpp>. A mathematically equivalent but longer representation of integrand is

```
{ adouble temp = arg;
          temp = temp*temp;
          val = sqrt(temp-1); }
```

The code block integrand may call on any elementary function that has already been defined in file <adolc/adouble.cpp>, so that one may also introduce iterated integrals.

9 Example Codes

The following listings are all simplified versions of codes that are contained in the example subdirectory <adolc-1.9.0>/examples of ADOL-C. In particular, we have left out timings, which are included in the complete codes.

9.1 Speelpenning's Example (speelpenning.cpp)

The first example evaluates the gradient and the Hessian of the function

$$y = f(x) = \prod_{i=0}^{n-1} x_i$$

using the appropriate drivers gradient and hessian.

```
#include <adolc/adouble.h>            // use of active doubles and taping
#include <adolc/drivers/drivers.h>    // use of "Easy to Use" drivers
                                      // gradient(.) and hessian(.)
```

```
...
void main() {
int n,i,j,tape_stats[11];
cout << "SPEELPENNINGS PRODUCT (ADOL-C Documented Example) \n";
cout << "number of independent variables = ?  \n";
cin >> n;
double* xp = new double[n];
double  yp = 0.0;
adouble* x = new adouble[n];
adouble  y = 1;
for(i=0;i<n;i++)
  xp[i] = (i+1.0)/(2.0+i);          // some initialization
trace_on(1);                        // tag =1, keep=0 by default
  for(i=0;i<n;i++) {
    x[i] <<= xp[i]; y *= x[i]; }
  y >>= yp;
  delete[] x;
trace_off();
tapestats(1,tape_stats);            // reading of tape statistics
cout<<"maxlive "<<tape_stats[2]<<"\n";
...                                 // ..... print other tape stats
double* g = new double[n];
gradient(1,n,xp,g);                 // gradient evaluation
double** H=(double**)malloc(n*sizeof(double*));
for(i=0;i<n;i++)
  H[i]=(double*)malloc((i+1)*sizeof(double));
hessian(1,n,xp,H);                  // H equals (n-1)g since g is
double errg = 0;                    // homogeneous of degree n-1.
double errh = 0;
for(i=0;i<n;i++)
  errg += fabs(g[i]-yp/xp[i]);      // vanishes analytically.
for(i=0;i<n;i++) {
  for(j=0;j<n;j++) {
    if (i>j)                        // lower half of hessian
      errh += fabs(H[i][j]-g[i]/xp[j]); } }
cout << yp-1/(1.0+n) << " error in function \n";
cout << errg <<" error in gradient \n";
cout << errh <<" consistency check \n";
}                                   // end main
```

9.2 Power Example (powexam.cpp)

The second example function evaluates the n-th power of a real variable x in $\log_2 n$ multiplications by recursive halving of the exponent. Since there is only one independent variable, the scalar derivative can be computed by using both forward and reverse, and the results are subsequently compared.

```
#include <adolc/adolc.h>           // use of ALL ADOL-C interfaces

adouble power(adouble x, int n) {
adouble z=1;
if (n>0) {                         // recursion and branches
  int nh =n/2;                     // that do not depend on
  z = power(x,nh);                 // adoubles are fine !!!!
  z *= z;
  if (2*nh != n)
    z *= x;
  return z; }                      // end if
else {
  if (n==0)                        // the local adouble z dies
    return z;                      // as it goes out of scope.
  else
    return 1/power(x,-n); }        // end else
} // end power
```

The function power above was obtained from the original undifferentiated version by simply changing the type of all doubles including the return variable to adoubles. The new version can now be called from within any active section, as in the following main program.

```
#include ...                       // as above
int main() {
int i,n,tag=1;
cout <<"COMPUTATION OF N-TH POWER (ADOL-C Documented Example)\n\n";
cout<<"monomial degree=? \n";      // input the desired degree
cin >> n;
                                   // allocations and initializations
double* Y[1];
*Y = new double[n+2];
double* X[1];                      // allocate passive variables with
*X = new double[n+4];              // extra dimension for derivatives
X[0][0] = 0.5;                     // function value = 0. coefficient
X[0][1] = 1.0;                     // first derivative = 1. coefficient
for(i=0;i<n+2;i++)
```

```
  X[0][i+2]=0;                    // further coefficients
double* Z[1];                     // used for checking consistency
*Z = new double[n+2];             // between forward and reverse
adouble y,x;                      // declare active variables
                                  // beginning of active section
trace_on(1);                      // tag = 1 and keep = 0
x <<= X[0][0];                    // only one independent var
y = power(x,n);                   // actual function call
y >>= Y[0][0];                    // only one dependent adouble
trace_off();                      // no global adouble has died
                                  // end of active section
double u[1];                      // weighting vector
u[0]=1;                           // for reverse call
for(i=0;i<n+2;i++) {              // note that keep = i+1 in call
  forward(tag,1,1,i,i+1,X,Y);     // evaluate the i-the derivative
  if (i==0)
    cout << Y[0][i] << " - " << y.value() << " = " << Y[0][i]-y.value()
    << " (should be 0)\n";
  else
    cout << Y[0][i] << " - " << Z[0][i] << " = " << Y[0][i]-Z[0][i]
    << " (should be 0)\n";
  reverse(tag,1,1,i,u,Z);         // evaluate the (i+1)-st derivative
  Z[0][i+1]=Z[0][i]/(i+1); }      // scale derivative to Taylorcoeff.
return 1;
}                                 // end main
```

Since this example has only one independent and one dependent variable, forward and reverse have the same complexity and calculate the same scalar derivatives, albeit with a slightly different scaling. By replacing the function power with any other univariate test function, one can check that forward and reverse are at least consistent. In the following example the number of independents is much larger than the number of dependents, which makes the reverse mode preferable.

9.3 Determinant Example (detexam.cpp)

Now let us consider an exponentially expensive calculation, namely, the evaluation of a determinant by recursive expansion along rows. The gradient of the determinant with respect to the matrix elements is simply the adjoint, i.e. the matrix of cofactors. Hence the correctness of the numerical result is easily checked by matrix-vector multiplication. The example illustrates the use of adouble arrays and pointers.

```
#include <adolc/adouble.h>        // use of active doubles and taping
```

```
#include <adolc/interfaces.h>      // use of basic forward/reverse
                                   // interfaces of ADOL-C
adouble** A;                       // A is an n x n matrix
int i,n;                           // k <= n is the order
adouble det(int k, int m) {        // of the sub-matrix
if (m == 0) return 1.0 ;           // its column indices
else {                             // are encoded in m
  adouble* pt = A[k-1];
  adouble t = zero;                // zero is predefined
  int s, p =1;
  if (k%2) s = 1; else s = -1;
  for(i=0;i<n;i++) {
    int p1 = 2*p;
    if (m%p1 >= p) {
      if (m == p) {
        if (s>0) t += *pt; else t -= *pt; }
      else {
        if (s>0)
          t += *pt*det(k-1,m-p);   // recursive call to det
        else
          t -= *pt*det(k-1,m-p); } // recursive call to det
      s = -s;}
    ++pt;
    p = p1;}
  return t; }
}                                  // end det
```

As one can see, the overloading mechanism has no problem with pointers and looks exactly the same as the original undifferentiated function except for the change of type from double to adouble. If the type of the temporary t or the pointer pt had not been changed, a compile time error would have resulted. Now consider a corresponding calling program.

```
#include ...                       // as above
int main() {
int i,j, m=1,tag=1,keep=1;
cout << "COMPUTATION OF DETERMINANTS (ADOL-C Documented Example)\n\n";
cout << "order of matrix = ? \n";  // select matrix size
cin >> n;
A = new adouble*[n];
trace_on(tag,keep);                // tag=1=keep
  double detout=0.0, diag = 1.0;   // here keep the intermediates for
  for(i=0;i<n;i++) {               // the subsequent call to reverse
```

```
    m *=2;
    A[i] = new adouble[n];          // not needed for adoublem
    adouble* pt = A[i];
    for(j=0;j<n;j++)
      A[i][j] <<= j/(1.0+i);        // make all elements of A independent
    diag += A[i][i].value();         // value() converts to double
    A[i][i] += 1.0; }
  det(n,m-1) >>= detout;            // actual function call
  printf("\n %f - %f = %f  (should be 0)\n",detout,diag,detout-diag);
trace_off();
double u[1];
u[0] = 1.0;
double* B = new double[n*n];
reverse(tag,1,n*n,1,u,B);
cout <<" \n first base? : ";
for (i=0;i<n;i++) {
  adouble sum = 0;
  for (j=0;j<n;j++)                 // the matrix A times the first n
    sum += A[i][j]*B[j];            // components of the gradient B
  cout<<sum.value()<<" "; }         // must be a Cartesian basis vector
return 1;
}                                   // end main
```

The variable diag should be mathematically equal to the determinant, because the matrix A is defined as a rank 1 perturbation of the identity.

9.4 Ordinary Differential Equation Example (odexam.cpp)

Here, we consider a nonlinear ordinary differential equation that is a slight modification of the Robertson test problem given in Hairer and Wanner's book on the numerical solution of ODEs [16]. The following source code computes the corresponding values of $y' \in \mathbb{R}^3$:

```
#include <adolc/adouble.h>            // use of active doubles and taping
#include <adolc/drivers/odedrivers.h> // use of "Easy To use" ODE drivers
#include <adolc/adalloc.h>            // use of ADOL-C allocation utilities

void tracerhs(short int tag, double* py, double* pyprime) {
adoublev y(3);                        // this time we left the parameters
adoublev yprime(3);                   // passive and use the vector types
trace_on(tag);
y <<= py;                             // initialize and mark independents
yprime[0] = -sin(y[2]) + 1e8*y[2]*(1-1/y[0]);
```

```
yprime[1] = -10*y[0] + 3e7*y[2]*(1-y[1]);
yprime[2] = -yprime[0] - yprime[1];
yprime >>= pyprime;                    // mark and pass dependents
trace_off(tag);
}                                      // end tracerhs
```

The Jacobian of the right-hand side has large negative eigenvalues, which make the ODE quite stiff. We have added some numerically benign transcendentals to make the differentiation more interesting. The following main program uses forode to calculate the Taylor series defined by the ODE at the given point y_0 and reverse as well as accode to compute the Jacobians of the coefficient vectors with respect to x_0.

```
#include .......                    // as above
int main() {
int i,j,deg;
int n=3;
double py[3];
double pyp[3];
cout << "MODIFIED ROBERTSON TEST PROBLEM (ADOL-C Documented Example)\n";
cout << "degree of Taylor series =?\n";
cin >> deg;
double **X;
X=(double**)malloc(n*sizeof(double*));
for(i=0;i<n;i++)
  X[i]=(double*)malloc((deg+1)*sizeof(double));
double*** Z=new double**[n];
double*** B=new double**[n];
short** nz = new short*[n];
for(i=0;i<n;i++) {
  Z[i]=new double*[n];
  B[i]=new double*[n];
  for(j=0;j<n;j++) {
    Z[i][j]=new double[deg];
    B[i][j]=new double[deg]; }     // end for
}                                  // end for
for(i=0;i<n;i++) {
  py[i] = (i == 0) ? 1.0 : 0.0;    // initialize the base point
  X[i][0] = py[i];                 // and the Taylor coefficient;
  nz[i] = new short[n]; }          // set up sparsity array
tracerhs(1,py,pyp);                // trace RHS with tag = 1
forode(1,n,deg,X);                 // compute deg coefficients
reverse(1,n,n,deg-1,Z,nz);         // U defaults to the identity
accode(n,deg-1,Z,B,nz);
```

```
cout << "nonzero pattern:\n";
for(i=0;i<n;i++) {
  for(j=0;j<n;j++)
    cout << nz[i][j]<<"\t";
  cout <<"\n"; }                    // end for
return 1;
}                                   // end main
```

The pattern nz returned by accode is

```
3  -1   4
1   2   2
3   2   4
```

The original pattern nz returned by reverse is the same except that the negative entry -1 was zero.

Acknowledgements

Parts of the ADOL-C source were developed by David Juedes, Hristo Mitev, Jay Srinivasan, Chuck Tyner, Olaf Vogel, and Duane Yoder. We are also indebted to George Corliss, Tom Epperly, Bruce Christianson, David Gay, Brad Karp, Koichi Kubota, Bob Olson, Marcela Rosemblun, and Dima Shiriaev for helping in various ways with the development and documentation of ADOL-C. Links to updated versions of the ADOL-C source code and this manual can be found on web page `http://www.math.tu-dresden.de/~adol-c`.

References

[1] Brett Averick, Jorge Moré, Christian Bischof, Alan Carle, and Andreas Griewank. Computing large sparse Jacobian matrices using automatic differentiation. Preprint MCS–P348–0193, Argonne National Laboratory, 1993.

[2] Martin Berz, Christian Bischof, George Corliss, and Andreas Griewank, editors. *Computational Differentiation – Techniques, Applications, and Tools.* Philadelphia, 1996. SIAM Proceedings of the Second International Workshop on Computational Differentiation, Santa Fe, New Mexico, February 12–14, 1996.

[3] Christian H. Bischof, Peyvand M. Khademi, Ali Bouaricha and Alan Carle. *Efficient computation of gradients and Jacobians by dynamic exploitation of sparsity in automatic differentiation*, Optimization Methods and Software, Vol.7(1):1-39, 1996.

[4] Kathryn E. Brenan, Stephen L. Campbell, and Linda R. Petzold. *Numerical Solution of Initial Value Problems in Differential-Algebraic Equations*. Elsevier (North Holland), 1989.

[5] Dan G. Cacuci. Sensitivity theory for nonlinear systems. I. Nonlinear functional analysis approach. *J. Math. Phys.*, 22(12):2794–2802, 1981.

[6] Dan G. Cacuci. Sensitivity theory for nonlinear systems. II. Extension to additional classes of responses. *J. Math. Phys.*, 22(12):2803–2812, 1981.

[7] Bruce Christianson. Reverse accumulation and accurate rounding error estimates for Taylor series. *Optimization Methods and Software*, 1:81–94, 1992.

[8] George Corliss, Christele Faure, Andreas Griewank, Laurent Hascoët, and Uwe Naumann, editors. *Automatic Differentiation: from Simulation to Optimization.* Springer Verlag, New York, 2001.

[9] Assefaw Gebremedhin, Fredrik Manne, and Alex Pothen. What color is your Jacobian? Graph coloring for computing derivatives. Technical report, Computer Science Dept, Old Dominion University, 2005. To appear in SIAM Review.

[10] Andreas Griewank: *Evaluating Derivatives, Principles and Techniques of Algorithmic Differentiation*. Frontiers in Appl. Math. 19, Phil., 2000.

[11] Andreas Griewank. Achieving logarithmic growth of temporal and spatial complexity in reverse automatic differentiation. *Optimization Methods and Software*, 1:35–54, 1992.

[12] Andreas Griewank and George F. Corliss, editors. *Differentiation of Algorithms: Theory, Implementation, and Application*. Philadelphia, 1991. SIAM Proceedings of the First International Workshop on Computational Differentiation, Breckenridge, Colorado, January 6-8, 1991.

[13] Andreas Griewank and Shawn Reese. On the calculation of Jacobian matrices by the Markowitz rule. In [12].

[14] Andreas Griewank, Jean Utke, and Andrea Walther. Evaluating higher derivative tensors by forward propagation of univariate Taylor series. *Mathematics of Computation*, 69:1117–1130, 2000.

[15] Andreas Griewank and Andrea Walther. Revolve: An Implementation of Checkpointing for the Reverse or Adjoint Mode of Computational Differentiation, *ACM Trans. Math. Software* 26:19–45, 2000.

[16] Ernst Hairer and Gerhard Wanner. *Solving Ordinary Differential Equations II.* Springer-Verlag, Berlin, 1991.

[17] Gershon Kedem. Automatic differentiation of computer programs. *ACM Trans. Math. Software*, 6(2):150–165, 1980.

[18] Donald E. Knuth. The Art of Computer Programming. Second edition. Addison-Wesley, Reading, 1973.

[19] Seppo Linnainmaa. Taylor expansion of the accumulated rounding error. *BIT (Nordisk Tidskrift for Informationsbehandling)*, 16(1):146–160, 1976.

[20] G. M. Ostrovskii, Yu. M. Volin, and W. W. Borisov. Über die Berechnung von Ableitungen. *Wissenschaftliche Zeitschrift der Technischen Hochschule für Chemie, Leuna-Merseburg*, 13(4):382–384, 1971.

[21] Louis B. Rall. *Automatic Differentiation: Techniques and Applications*, Volume 120 of *Lecture Notes in Computer Science*. Springer Verlag, Berlin, 1981.

[22] Louis B. Rall. Differentiation and generation of Taylor coefficients in Pascal-SC. In Ulrich W. Kulisch and Willard L. Miranker, editors, *A New Approach to Scientific Computation*, pages 291–309. Academic Press, New York, 1983.

[23] Bert Speelpenning. *Compiling Fast Partial Derivatives of Functions Given by Algorithms*. Ph.D. thesis, Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana-Champaign, Ill., January 1980.

[24] Olivier Talagrand and Philippe Courtier. Variational assimilation of meteorological observations with the adjoint vorticity equation – Part I. Theory. *Q. J. R. Meteorol. Soc.*, 113:1311–1328, 1987.

[25] Andrea Walther. Computing Sparse Hessians with Automatic Differentiation. Preprint MATH-WR-03-2005, TU Dresden, 2005.

[26] R. E. Wengert. A simple automatic derivative evaluation program. *Comm. ACM*, 7(8):463–464, 1964.

[27] Paul Werbos. *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*. Ph.D. thesis, Committee on Applied Mathematics, Harvard University, Cambridge, Mass., November 1974.