

AD *in* Fortran

Part 2: Implementation *via* Preprocessor

Alexey Radul, Barak A. Pearlmutter, and Jeffrey Mark Siskind

Abstract We describe an implementation of the FARFEL FORTRAN AD extensions (Radul et al., 2012). These extensions integrate forward and reverse AD directly into the programming model, with attendant benefits to flexibility, modularity, and ease of use. The implementation we describe is a “preprocessor” that generates input to existing FORTRAN-based AD tools. In essence, blocks of code which are targeted for AD by FARFEL constructs are put into subprograms which capture their lexical variable context, and these are closure-converted into top-level subprograms and specialized to eliminate **EXTERNAL** arguments, rendering them amenable to existing AD preprocessors, which are then invoked, possibly repeatedly if the AD is nested.

Key words: Nesting, multiple transformation, forward mode, reverse mode, TAPE-NADE, ADIFOR, programming-language implementation

1 Introduction

The *Forward And Reverse Fortran Extension Language* (FARFEL) extensions to FORTRAN enable smooth and modular use of AD (Radul et al., 2012). A variety of implementation strategies present themselves, ranging from (a) deep integration into a FORTRAN compiler, to (b) a preprocessor that performs the requested AD and generates FORTRAN or some other high-level language, to (c) a preprocessor which itself does no AD but generates input to an existing FORTRAN AD preprocessor. This last strategy leverages existing FORTRAN-based AD tools and compilers,

Alexey Radul
Hamilton Inst., National Univ. Ireland Maynooth, alexey.radul@nuim.ie

Barak A. Pearlmutter
Dept. Computer Sci. & Hamilton Inst., National Univ. Ireland Maynooth, barak@cs.nuim.ie

Jeffrey Mark Siskind
Electrical and Computer Engineering, Purdue University, IN, USA, qobi@purdue.edu

avoiding re-implementation of the AD transformations, at the expense of inheriting some of the limitations of the AD tool it invokes. This is the technique we will focus upon.

FARFALLEN transforms FARFEL input into FORTRAN, and invokes an existing AD system (Bischof et al., 1992; Hascoët and Pascual, 2004) to generate the needed derivatives. The process can make use of a variety of existing FORTRAN-based AD preprocessors, making it easy for the programmer to switch between them. There is significant semantic mismatch between FARFEL and the AD operations allowed by the AD systems used, necessitating rather dramatic code transformations. When the FARFEL program involves nested AD, the transformations and staging become even more involved. Viewed as a whole, this tool automates the task of applying AD, including the detailed maneuvers required for nested application of existing tools, thereby extending the reach and utility of AD.

The remainder of the paper is organized as follows: Section 2 reviews the FARFEL extensions. A complete example program on page 4 illustrates their use. Section 3 describes the implementation in detail using this example program. Section 4 summarizes work’s contributions.

2 Language Extensions

FARFEL provides two principal extensions to FORTRAN: syntax for AD and for nested subprograms.

Extension 1: AD Syntax

FARFEL adds the **ADF** construct for forward AD:

```
ADF (TANGENT (var) = expr ...)
  statements
END ADF (var = TANGENT (var) ...)
```

Multiple opening and closing assignments are separated by commas. Independent variables are listed in the “calls” to **TANGENT** on the left-hand sides of the opening assignments and are given the specified tangent values. Dependent variables appear in the “calls” to **TANGENT** on the right-hand sides of the closing assignments and the corresponding tangent values are assigned to the indicated destination variables. The **ADF** construct uses forward AD to compute the directional derivative of the dependent variables at the point specified by the vector of independent variables in the direction specified by the vector of tangent values for the independent variables and assigns it to the destination variables.

An analogous FARFEL construct supports reverse AD:

```
ADR (COTANGENT (var) = expr ...)
  statements
END ADR (var = COTANGENT (var) ...)
```

Dependent variables are listed in the “calls” to **COTANGENT** on left-hand sides of the opening assignments and are given the specified cotangent values as inputs to

the reverse phase. *Independent* variables appear in the “calls” to `COTANGENT` on the right-hand sides of the closing assignments and the corresponding cotangent values at the end of the reverse phase are assigned to the indicated destination variables. The `ADR` construct uses reverse AD to compute the gradient with respect to the independent variables at the point specified by the vector of independent variables induced by the specified gradient with respect to the dependent variables, and assigns it to the destination variables. The expressions used to initialize the cotangent inputs to the reverse phase are evaluated at the end of the forward phase, even though they appear textually prior to the statements specifying the forward phase. This way, the direction input to the reverse phase can depend on the result of the forward phase.

For both `ADF` and `ADR`, implied-`DO` syntax is used to allow arrays in the opening and closing assignments. By special dispensation, the statement `ADF(var)` is interpreted as `ADF(TANGENT(var)=1)` and `ADR(var)` as `ADR(COTANGENT(var)=1)`.

Extension 2: Nested Subprograms

In order to conveniently support distinctions between different variables of differentiation for distinct invocations of AD, as in the example below, we borrow from ALGOL 60 (Backus et al., 1963) and generalize the FORTRAN “statement function” construct by allowing subprograms to be defined inside other subprograms, with lexical scope. As in ALGOL 60, the scope of parameters and declared variables is the local subprogram, and these may shadow identifiers from the surrounding scope. Implicitly declared variables have the top-level subprogram as their scope.

Concrete Example

In order to describe the implementation of the above constructs, we employ a concrete example. The task is to find an equilibrium (a^*, b^*) of a two-player game with continuous scalar strategies a and b and given payoff functions A and B . The method is to find roots of

$$a^* = \underset{a}{\operatorname{argmax}} A(a, \underset{b}{\operatorname{argmax}} B(a^*, b)) \quad (1)$$

The full program is given, for reference, in Listing 1. The heart of the program is the implementation `EQLBRM` of (1). Note that this whole program is only 63 lines of code, with plenty of modularity boundaries. This code is used as a running example for the remainder of the paper.

3 Implementation

FARFEL is implemented by the FARFALLEN preprocessor. The current version is merely a proof of concept, and not production quality: it does not accept the entire FORTRAN77 language, and does not scale. However, its principles of operation will be unchanged in a forthcoming production-quality implementation. Here we describe the reduction of FARFEL constructs to FORTRAN, relying on existing FORTRAN-based AD tools for the actual derivative transformations.

Listing 1 Complete example FARFEL program: equilibria of a continuous-strategy game.

```

C   ASTAR & BSTAR: GUESSES IN, OPTIMIZED VALUES OUT
SUBROUTINE EQLBRM(BIGA, BIGB, ASTAR, BSTAR, N)
EXTERNAL BIGA, BIGB
  FUNCTION F(ASTAR)
    FUNCTION G(A)
      FUNCTION H(B)
        H = BIGB(ASTAR, B)
      END
    BSTAR = ARGMAX(H, BSTAR, N)
    G = BIGA(A, BSTAR)
  END
  F = ARGMAX(G, ASTAR, N)-ASTAR
END
ASTAR = ROOT(F, ASTAR, N)
END

FUNCTION ROOT(F, X0, N)
X = X0
DO 1669 I=1,N
CALL DERIV2(F, X, Y, YPRIME)
1669 X = X-Y/YPRIME
ROOT = X
END

SUBROUTINE DERIV2(F, X, Y, YPRIME)
EXTERNAL F
ADF(X)
Y = F(X)
END ADF(YPRIME = TANGENT(Y))
END

FUNCTION ARGMAX(F, X0, N)
  FUNCTION FPRIME(X)
    FPRIME = DERIV1(F, X)
  END
  ARGMAX = ROOT(FPRIME, X0, N)
END

FUNCTION DERIV1(F, X)
EXTERNAL F
ADF(X)
Y = F(X)
END ADF(DERIV1 = TANGENT(Y))
END

FUNCTION GMBIGA(A, B)
PRICE = 20-0.1*A-0.1*B
COSTS = A*(10-0.05*A)
GMBIGA = A*PRICE-COSTS
END

FUNCTION GMBIGB(A, B)
PRICE = 20-0.1*B-0.0999*A
COSTS = B*(10.005-0.05*B)
GMBIGB = B*PRICE-COSTS
END

PROGRAM MAIN
READ *, ASTAR
READ *, BSTAR
READ *, N
CALL EQLBRM(GMBIGA, GMBIGB, ASTAR, BSTAR, N)
PRINT *, ASTAR, BSTAR
END

```

FARFEL introduces two new constructs into FORTRAN: nested subprograms and syntax for requesting AD. We implement nested subprograms by extracting them to the top level, and communicating the free variables from the enclosing subprogram by passing them as arguments into the new top-level subprogram. This is an instance of *closure conversion*, a standard class of techniques for converting nested subprograms to top-level ones (Johnsson, 1985). In order to accommodate passing formerly-free variables as arguments, we must adjust all the call sites of the formerly-nested subprogram; we must specialize all the subprograms that accept that subprogram as an external to also accept the extra closure parameters; and adjust all call sites to all those specialized subprograms to pass those extra parameters.

We implement the AD syntax by constructing new subroutines that correspond to the statements inside each **ADF** or **ADR** block, arranging to call the AD tool of choice on each of those new subroutines, and transforming the block itself into a call to the appropriate tool-generated subroutine.

Nested Subprograms in Detail

Let us illustrate closure conversion on our example. Recall `ARGMAX` in our example program:

```

FUNCTION ARGMAX(F, X0, N)
  FUNCTION FPRIME(X)
    FPRIME = DERIV1(F, X)
  END
  ARGMAX = ROOT(FPRIME, X0, N)
END

```

This contains the nested function `FPRIME`. We closure convert this as follows. First, extract `FPRIME` to the top level:

```

FUNCTION ARGMAX_FPRIME(X, F)
  ARGMAX_FPRIME = DERIV1(F, X)
END

FUNCTION ARGMAX(F, X0, N)
  ARGMAX = ROOT_1(ARGMAX_FPRIME, F, X0, N)
END

```

Note the addition of a closure argument for `F` since it is freely referenced in `FPRIME`, and the addition of the same closure argument at the call site, since `FPRIME` is passed as an external to `ROOT`. Then we specialize `ROOT` to create a version that accepts the needed set of closure arguments (in this case one):

```

FUNCTION ROOT_1(F, F1, X0, N)
  X = X0
  DO 1669 I=1,N
    CALL DERIV2_1(F, F1, X, Y, YPRIME)
  1669 X = X-Y/YPRIME
  ROOT_1 = X
END

```

Since `ROOT` contained a call to `DERIV2`, passing it the external passed to `ROOT`, we must also specialize `DERIV2`:

```

SUBROUTINE DERIV2_1(F, F1, X, Y, YPRIME)

```

```

EXTERNAL F
ADF(X)
Y = F(X, F1)
END ADF(YPRIME = TANGENT(Y))
END

```

We must, in general, copy and specialize the portion of the call graph where the nested subprogram travels, which in this case is just two subprograms. During such copying and specialization, we propagate external constants (e.g., `FPRIME` through the call to `ROOT_1`, the call to `DERIV2_1`, and the call site therein) allowing the elimination of the `EXTERNAL` declaration for these values. This supports AD tools that do not allow taking derivatives through calls to external subprograms.

That is the process for handling one nested subprogram. In our example, the same is done for `F` in `EQLBRM`, `G` in `F`, and `H` in `G`. Doing so causes the introduction of a number of closure arguments, and the specialization of a number of subprograms to accept those arguments; including perhaps further specializing things that have already been specialized. The copying also allows a limited form of subprogram reentrancy: even if recursion is disallowed (as in traditional FORTRAN77) our nested uses of `ARGMAX` will cause no difficulties because they will end up calling two different specializations of `ARGMAX`.

Note that we must take care to prevent this process from introducing spurious aliases. For example, in `EQLBRM`, the internal function `F` that is passed to `ROOT` closes over the iteration count `N`, which is also passed to `ROOT` separately. When specializing `ROOT` to accept the closure parameters of `F`, we must not pass `N` to the specialization of `ROOT` twice, lest we run afoul of FORTRAN's prohibition against assigning to aliased values. Fortunately, such situations are syntactically apparent.

Finally, specialization leaves behind unspecialized (or underspecialized) versions of subprograms, which may now be unused, and if so can be eliminated. In this case, that includes `ROOT`, `DERIV2`, `ARGMAX`, `FPRIME`, and `DERIV1` from the original program, as well as some intermediate specializations thereof.

AD Syntax in Detail

We implement the AD syntax by first canonicalizing each `ADF` or `ADR` block to be a single call to a (new, internal) subroutine, then extracting those subroutines to the top level, then rewriting the block to be a call to an AD-transformed version of the subroutine, and then arranging to call the AD tool of choice on each of those new subroutines to generate the needed derivatives.

Returning to our example program, closure conversion of nested subprograms produced the following specialization of `DERIV1`:

```

FUNCTION DERIV1_1(ASTAR, BIGA, BIGB, BSTAR, N, X)
ADF(X)
Y = EQLBRM_F_G(X, ASTAR, BIGA, BIGB, BSTAR, N)
END ADF(DERIV1_1 = TANGENT(Y))
END

```

which contains an `ADF` block. We seek to convert this into a form suitable for invoking the AD preprocessor. We first canonicalize by introducing a new subroutine to capture the statements in the `ADF` block, producing the following:

```

FUNCTION DERIV1_1(ASTAR, BIGA, BIGB, BSTAR, N, X)
  SUBROUTINE ADF1()
    Y = EQLBRM_F_G(X, ASTAR, BIGA, BIGB, BSTAR, N)
  END
  ADF(X)
  CALL ADF1()
END ADF(DERIV1_1 = TANGENT(Y))
END

```

Extracting the subroutine ADF1 to the top level as before yields the following:

```

SUBROUTINE DERIV1_1_ADF1(X, ASTAR, BIGA, BIGB, BSTAR, N, Y)
  Y = G(X, ASTAR, BIGA, BIGB, BSTAR, N)
END

FUNCTION DERIV1_1(ASTAR, BIGA, BIGB, BSTAR, N, X)
  ADF(X)
  CALL DERIV1_1_ADF1(X, ASTAR, BIGA, BIGB, BSTAR, N, Y)
END ADF(DERIV1_1 = TANGENT(Y))
END

```

Now we are properly set up to rewrite the **ADF** block into a subroutine call—specifically, to a subroutine that will be generated from DERIV1_1_ADF1 by AD. The exact result depends on the AD tool that will be used to construct the derivative of DERIV1_1_ADF1; for TAPENADE, the generated code looks like this:

```

FUNCTION DERIV1_1(ASTAR, BIGA, BIGB, BSTAR, N, X)
  X_G1 = 1
  ASTAR_G1 = 0
  BSTAR_G1 = 0
  CALL DERIV1_1_ADF1_G1(X, X_G1, ASTAR, ASTAR_G1, BIGA, BIGB,
+BSTAR, BSTAR_G1, N, Y, DERIV1_1)
END

```

Different naming conventions are used for DERIV1_1_ADF1_G1 when generating code for ADIFOR; the parameter passing conventions of TAPENADE and ADIFOR agree in this case. FARFALLEN maintains the types of variables in order to know whether to generate variables to hold tangents and cotangents (which are initialized to zero if they were not declared in the relevant opening assignments.)

The same must be repeated for each **ADF** and **ADR** block; in our example there are five in all: two in specializations of DERIV1 and three in specializations of DERIV2. We must also specialize EQLBRM and its descendants in the call graph, by the process already illustrated, to remove external calls to the objective functions, for the reasons described earlier.

Finally, we must invoke the user's preferred AD tool to generate all the needed derivatives. Here, FARFALLEN might invoke TAPENADE as follows:

```

#!/bin/sh
tapenade -root deriv1_2_adf2 -d -o eqlbrm42 -diffvarname "_g2"\
-difffuncname "_g2" eqlbrm42.f
tapenade -root deriv2_1_2_adf4 -d -o eqlbrm42 -diffvarname "_g4"\
-difffuncname "_g4" eqlbrm42{,_g2}.f
tapenade -root deriv1_1_adf1 -d -o eqlbrm42 -diffvarname "_g1"\
-difffuncname "_g1" eqlbrm42{,_g2,_g4}.f
tapenade -root deriv2_1_1_adf3 -d -o eqlbrm42 -diffvarname "_g3"\
-difffuncname "_g3" eqlbrm42{,_g2,_g4,_g1}.f

```

```
tapenade -root deriv2_2_adf5 -d -o eqlbrm42 -diffvarname "_g5" \
-difffuncname "_g5" eqlbrm42{,_g2,_g4,_g1,_g3}.f
```

We must take care that multiple invocations of the AD tool to generate the various derivatives occur in the proper order, which is computable from the call graph of the program, to ensure that generated derivative codes that are used in subprograms to be differentiated are available to be transformed. For example, all the derivatives of `EQLBRM_F_G_H` needed for its optimizations must be generated before generating derivatives of (generated subprograms that call) `EQLBRM_F_G`. We must also take care to invoke the AD tool with different prefixes/suffixes, so that variables and subprograms created by one differentiation do not clash with those made by another.

Performance

We tested the performance of code generated by FARFALLEN in conjunction with TAPENADE and with ADIFOR. We executed FARFALLEN once to generate FORTRAN77 source using TAPENADE for the automatic differentiation, and separately using ADIFOR for the AD. In each case, we compiled the resulting program (`gfortran 4.6.2-9`, 64-bit Debian sid, `-Ofast -fwhole-program`, single precision) with $N = 1000$ iterations at each level and timed the execution of the binary on a 2.93GHz Intel i7 870. For comparison, we translated the same computation into VLAD (Pearlmutter and Siskind, 2008) (Fig. 1), compiled it with STALINGRAD (Siskind and Pearlmutter, 2008), and ran it on the same machine. STALINGRAD has the perhaps unfair advantage of being an optimizing compiler with integrated support for AD, so we are pleased that FARFALLEN was able to achieve performance that was nearly competitive.

TAPENADE	ADIFOR	STALINGRAD
6.97	8.92	5.83

The VLAD code in Fig. 1 was written with the same organization, variable names, subprogram names, and parameter names and order as the corresponding FARFEL code in Listing 1 to help a reader unfamiliar with SCHEME, the language on which VLAD is based, understand how FARFEL and VLAD both represent the same essential notions of nested subprograms and the AD discipline of requesting derivatives precisely where they are needed. Because functional-programming languages, like SCHEME and VLAD, prefer higher-order functions (i.e., operators) over the block-based style that is prevalent in FORTRAN, AD is invoked via the \vec{J} and \overleftarrow{J} operators rather than via the `ADF` and `ADR` statements. However, there is a one-to-one correspondence between

$$\vec{J} : f \times x \times \acute{x} \rightarrow y \times \acute{y}$$

an operator that takes a function f as input, along with a primal argument x and tangent argument \acute{x} and returns both a primal result y and a tangent result \acute{y} , and the following:

```
ADF(TANGENT(x) =  $\acute{x}$ )
y = f(x)
END ADF( $\acute{y}$  = TANGENT(y))
```



```

EQLBRM(A,B,a0,b0,n)  $\triangleq$  let f(a')  $\triangleq$  let g(a)  $\triangleq$  let h(b)  $\triangleq$  B(a',b)
                                in A(a, ARGMAX(h,b0,n))
                                in ARGMAX(g,a',n) - a'
                                a* = ROOT(f,a0,n)
                                in let h(b)  $\triangleq$  B(a*,b)
                                    b* = ARGMAX(h,b0,n)
                                in a*,b*

ROOT(f,x,n)  $\triangleq$  if n = 0
                then x
                else let y,y' = DERIV2(f,x)
                    in ROOT(f,x -  $\frac{y}{y'}$ , n - 1)

DERIV2(f,x)  $\triangleq$   $\vec{J}$ (f,x,1)
ARGMAX(f,x0,n)  $\triangleq$  let f'(x)  $\triangleq$  DERIV1(f,x)
                in ROOT(f',x0,n)
DERIV1(f,x)  $\triangleq$  let x,y =  $\vec{J}$ (f,x,1)
                in y

A(a,b)  $\triangleq$  let price = 20 - 0.1  $\times$  a - 0.1  $\times$  b
          costs = a  $\times$  (10 - 0.05  $\times$  a)
          in a  $\times$  price - costs

B(a,b)  $\triangleq$  let price = 20 - 0.1  $\times$  b - 0.0999  $\times$  a
          costs = b  $\times$  (10.005 - 0.05  $\times$  b)
          in b  $\times$  price - costs

let a0 = READREAL()
    b0 = READREAL()
    n = READREAL()
    a*,b* = EQLBRM(A,B,a0,b0,n)
in WRITEREAL(a*)
    WRITEREAL(b*)

```

Fig. 1 Complete VLAD program for our concrete example with the same organization and functionality as the FARFEL program in Listing 1.

Similarly, there is a one-to-one correspondence between

$$\overleftarrow{\mathbf{J}} : f \times x \times \dot{y} \rightarrow y \times \dot{x}$$

an operator that takes a function f as input, along with a primal argument x and cotangent \dot{y} and returns both a primal result y and a cotangent \dot{x} , and

```

ADR(COTANGENT(y) =  $\dot{y}$ )
y = f(x)
END ADR( $\dot{x}$  = COTANGENT(x))

```

The strong analogy between how the callee-derives AD discipline is represented in both FARFEL and VLAD serves two purposes: it enables the use of the STALINGRAD

compiler technology for compiling FARFEL and facilitates migration from legacy imperative languages to more modern, and more easily optimized, pure languages.

4 Conclusion

We have illustrated an implementation of the FARFEL extensions to FORTRAN—nested subprograms and syntax for AD (Radul et al., 2012). These extensions enable convenient, modular programming using a callee-derives paradigm of automatic differentiation. Our implementation is a preprocessor that translates FARFEL FORTRAN extensions into input suitable for an existing AD tool. This strategy enables modular, flexible use of AD in the context of an existing legacy language and tool chain, without sacrificing the desirable performance characteristics of these tools: in the concrete example, only 20%–50% slower than a dedicated AD-enabled compiler, depending on which FORTRAN AD system is used.

Acknowledgements This work was supported, in part, by Science Foundation Ireland grant 09/IN.1/12637, National Science Foundation grant CCF-0438806, the Naval Research Laboratory under Contract Number N00173-10-1-G023, and the Army Research Laboratory accomplished under Cooperative Agreement Number W911NF-10-2-0060. Any views, opinions, findings, conclusions, or recommendations contained or expressed in this document or material are those of the author(s) and do not necessarily reflect or represent the views or official policies, either expressed or implied, of SFI, NSF, NRL, the Office of Naval Research, ARL, or the Irish or U.S. Governments. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes, notwithstanding any copyright notation herein.

References

- J. W. Backus, F. L. Bauer, J. Green, C. Katz, J. McCarthy, P. Naur, A. J. Perlis, H. Rutishauser, K. Samelson, B. Vauquois, J. H. Wegstein, A. van Wijngaarden, and M. Woodger. Revised report on the algorithmic language ALGOL 60. *The Computer Journal*, 5(4):349–367, 1963. doi: 10.1093/comjnl/5.4.349.
- Christian H. Bischof, Alan Carle, George F. Corliss, Andreas Griewank, and Paul D. Hovland. ADIFOR: Generating derivative codes from Fortran programs. *Scientific Programming*, 1(1):11–29, 1992.
- Laurent Hascoët and Valérie Pascual. TAPENADE 2.1 user’s guide. Rapport technique 300, INRIA, Sophia Antipolis, 2004. URL <http://www.inria.fr/rrrt/rt-0300.html>.
- Thomas Johnsson. Lambda lifting: Transforming programs to recursive equations. In *Functional Programming Languages and Computer Architecture*, Nancy, France, September 1985. Springer Verlag.
- Barak A. Pearlmutter and Jeffrey Mark Siskind. Using programming language theory to make automatic differentiation sound and efficient. In Christian H.

Bischof, H. Martin Bückner, Paul D. Hovland, Uwe Naumann, and J. Utke, editors, *Advances in Automatic Differentiation*, volume 64 of *Lecture Notes in Computational Science and Engineering*, pages 79–90. Springer, Berlin, 2008. ISBN 978-3-540-68935-5. doi: 10.1007/978-3-540-68942-3_8.

Alexey Radul, Barak A. Pearlmutter, and Jeffrey Mark Siskind. AD in Fortran, Part 1: Design, 2012. URL <http://arxiv.org/abs/1203.1448>.

Jeffrey Mark Siskind and Barak A. Pearlmutter. Using polyvariant union-free flow analysis to compile a higher-order functional-programming language with a first-class derivative operator to efficient Fortran-like code. Technical Report TR-ECE-08-01, School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA, jan 2008. URL <http://docs.lib.purdue.edu/ecetr/367>.