

*On two developments of the adjoint model of Algorithmic
Differentiation to meet the needs of computationally
intensive applications*

A. TAFTAF

Thesis Defense Presentation

17 January 2017

Algorithmic Differentiation (AD)

- Derivatives of functions are required in many areas of computational science, e.g. design optimization.
- For complex function $F : X \in \mathbf{R}^N \rightarrow Y \in \mathbf{R}^M$, expressed as a computer program P , a method known as **Algorithmic Differentiation** produces code that computes the derivatives of this functions.
- AD considers the original program P as a set of successive instructions $\{l_1; l_2; \dots; l_p\}$ that computes a set of elementary differentiable functions $\{g_1, g_2, \dots, g_p\}$, each by instruction l_k , so that:

$$F(X) = g_p(\dots(g_2(g_1(X))\dots)) \quad (1)$$

Tangent and Adjoint modes of AD

- Applying the chain rule to $F(X)$ gives a new function $F'(X)$ that computes the first order full derivatives, i.e. the Jacobian:

$$F'(X) = g'_p(X_{p-1}) * \dots * g'_2(X_1) * g'_1(X) \quad (2)$$

- We may extend every instruction l_k , so that it computes in addition to the elementary function g_k , its derivative g'_k .
- Computing the whole Jacobian $F'(X)$ may require too much time and memory.
- In practice, we compute one of these two projections:

$$F'(X) * \dot{X} \quad \text{or} \quad \bar{Y} * F'(X)$$

where \dot{X} is a vector in R^N and \bar{Y} is a row-vector in R^M .

- The particular modes of AD that compute these projections are called respectively the **tangent and adjoint modes**.

Adjoint mode of AD

- When the number of inputs is much larger than the number of outputs, e.g. gradients, the adjoint mode is recommended.
- The adjoint mode computes $\bar{X} = \bar{Y} * F'(X)$. Recalling equation 2, we get:

$$\bar{X} = \bar{Y} * g'_p(X_{p-1}) * \dots * g'_2(X_1) * g'_1(X) \quad (3)$$

- Equation 3 is best computed from left to right.
- Primal values are needed in the opposite of their computation order, i.e. X_{p-1} is needed first, then X_{p-2} , then X_{p-3} , etc.

Dealing with Data-Flow reversal of Adjoint mode

In the Store-All context, the adjoint code consist of two sweeps:

- First sweep (the forward sweep) runs the original program and saves the intermediate values.
- Second sweep (the backward sweep) computes the partial derivatives, using the intermediate primal values already stored during the first sweep..

```
SUBROUTINE FOO(x, y )
```

```
  w = x * x
```

```
  w = 2 * w3
```

```
  w = sin(w)
```

```
  y = w
```

```
END
```

Dealing with Data-Flow reversal of Adjoint mode

SUBROUTINE $\overline{FOO}(x, \bar{x}, y, \bar{y})$

push(w)

$w = x * x$

push(w)

$w = 2 * w^3$

push(w)

$w = \sin(w)$

push(y)

$y = w$

pop(y)

$\bar{w} = \bar{y}; \bar{y} = 0$

pop(w)

$\bar{w} = \bar{w} * \cos(w)$

pop(w)

$\bar{w} = 6 * \bar{w} * w^2$

pop(w)

$\bar{x} = 2 * \bar{w} * x$

$\bar{w} = 0$

END

Motivations

- This thesis is part of AboutFlow project which focuses on methods of optimization gradient-based.
- Adjoint mode of Algorithmic Differentiation (AD) is particularly attractive for computing gradients
- Adjoint mode requires data-flow reversal which is very expensive.

⇒ How to reduce the cost of Adjoint mode of AD ?

Contents

- 1 Introduction
- 2 An efficient Adjoint of Fixed-Point Loops
- 3 Summary and Further Work

Fixed-Point (FP) loops

- Many equations $F(z, x) = 0$ may be solved by using iterative methods:

$z = \text{initial_guess}$

Do while z not converged

$z = \phi(z, x)$

end

$y = f(z, x)$

- state : variable that reaches a stationary value: z
- parameters : variables that influence the result but never modified: x

Standard adjoint of Fixed-Point loops: Black Box

`z = initial_guess`

Do while z not converged

`push(z)`

`z = $\phi(z, x)$`

end

`y = f(z, x)`

$$\bar{x} = \bar{y} * \frac{\partial f}{\partial \mathbf{x}}(z, x)$$

$$\bar{z} = \bar{y} * \frac{\partial f}{\partial \mathbf{z}}(z, x)$$

Do "as many times"

`pop(z)`

$$\bar{x} += \bar{z} * \frac{\partial \phi}{\partial \mathbf{x}}(z, x)$$

$$\bar{z} = \bar{z} * \frac{\partial \phi}{\partial \mathbf{z}}(z, x)$$

end

Efficient adjoint of Fixed-Point loops: Piggyback

- Many FP loops satisfy a FP equation of the form:
 $z = z - P * F(z, x)$.
- The adjoint of these loops may be written as:

Repeat until \bar{w} and \bar{z} stables

$$w = F(z, x)$$

$$\bar{z} = \bar{w} * \frac{\partial F}{\partial z}(z, x) + \bar{y} * \frac{\partial f}{\partial z}(z, x)$$

$$z = z - P * w$$

$$\bar{w} = \bar{w} - \bar{z} * P$$

end

$$\bar{x} = \bar{w} * \frac{\partial F}{\partial x}(z, x) + \bar{y} * \frac{\partial f}{\partial x}(z, x)$$

Refinements of the Piggyback approach

- **Delayed Piggyback:** Consists in applying Piggyback after that the original Fixed-Point loop has “sufficiently” converged to the solution.
- **Blurred Piggyback:** Piggyback can be applied in a parallel way: a process that computes the original values and another process that computes the adjoint values. The Blurred Piggyback proposes to run these processes in an asynchronous way.

Efficient adjoint of Fixed-Point loops: Two-Phases

Do while z not converged

$$z = \phi(z, x)$$

end

$$z = \phi(z, x)$$

store(z)

$$y = f(z, x)$$

$$\bar{x} = \bar{y} * \frac{\partial f}{\partial x}(z, x)$$

$$\bar{z} = \bar{y} * \frac{\partial f}{\partial z}(z, x)$$

$$\bar{w} = \bar{z}; \bar{x}_0 = \bar{x}$$

Do while \bar{w} not converged

restore(z)

$$\bar{w} = \bar{w} * \frac{\partial \phi}{\partial z}(z, x) + \bar{z}$$

$$\bar{x} = \bar{w} * \frac{\partial \phi}{\partial x}(z, x) + \bar{x}_0$$

end

Refinement of the Two-Phases approach

Do while z not converged

$$z = \phi(z, x)$$

end

$$z = \phi(z, x)$$

store(z)

$$y = f(z, x)$$

$$\bar{x} = \bar{y} * \frac{\partial f}{\partial x}(z, x)$$

$$\bar{z} = \bar{y} * \frac{\partial f}{\partial z}(z, x)$$

$$\bar{w} = \bar{z}; \bar{x}_0 = \bar{x}$$

Do while \bar{w} not converged

restore(z)

$$\bar{w} = \bar{w} * \frac{\partial \phi}{\partial z}(z, x) + \bar{z}$$

end

$$\bar{x} = \bar{w} * \frac{\partial \phi}{\partial x}(z, x) + \bar{x}_0$$

Refinement of the Black Box approach

Do while z not converged

$$z = \phi(z, x)$$

end

$$z = \phi(z, x)$$

store(z)

$$y = f(z, x)$$

$$\bar{x} = \bar{y} * \frac{\partial f}{\partial x}(z, x)$$

$$\bar{z} = \bar{y} * \frac{\partial f}{\partial z}(z, x)$$

Do "as many times"

restore(z)

$$\bar{x} += \bar{z} * \frac{\partial \phi}{\partial x}(z, x)$$

$$\bar{z} = \bar{z} * \frac{\partial \phi}{\partial z}(z, x)$$

end

Our choices

We prefer the approach that:

- Does not make assumptions on the shape of the FP loop and easy to implement. This is not the case of Piggyback class of methods.
- Considers that the adjoint of the FP loop is a FP loop itself. This is not the case of Black Box class of methods.
- Computes \bar{x} outside the adjoint loop. This is not the case of the Two-Phases approach.

⇒ We select the **Refined Two-Phases approach** to be implemented in our AD tool Tapenade.

Automatic Detection of Fixed-Point elements

- One directive to identify FP loop.
- State and parameters detected automatically from static data-flow analysis :

$$\begin{aligned} \text{state} &= \mathbf{out}(\text{FP loop}) \cap \mathbf{live} \\ \text{parameters} &= \mathbf{use}(\text{FP loop}) \setminus \mathbf{out}(\text{FP loop}) \end{aligned} \quad (4)$$

use : differentiable variables read by the FP loop

out : differentiable variables written by the FP loop

live : differentiable variables used in the sequel of FP loop

Implementation: Enable a repeated access to the stack

Do while z not converged

$$z = \phi(z, x)$$

end

$$z = \phi(z, x)$$

push(z)

$$y = f(z, x)$$

$$\bar{x} = \bar{y} * \frac{\partial f}{\partial x}(z, x)$$

$$\bar{z} = \bar{y} * \frac{\partial f}{\partial z}(z, x)$$

$$\bar{w} = \bar{z}$$

CALL start_repeat_stack()

Do while \bar{w} not converged

pop(z)

$$\bar{w} = \bar{w} * \frac{\partial \phi}{\partial z}(z, x) + \bar{z}$$

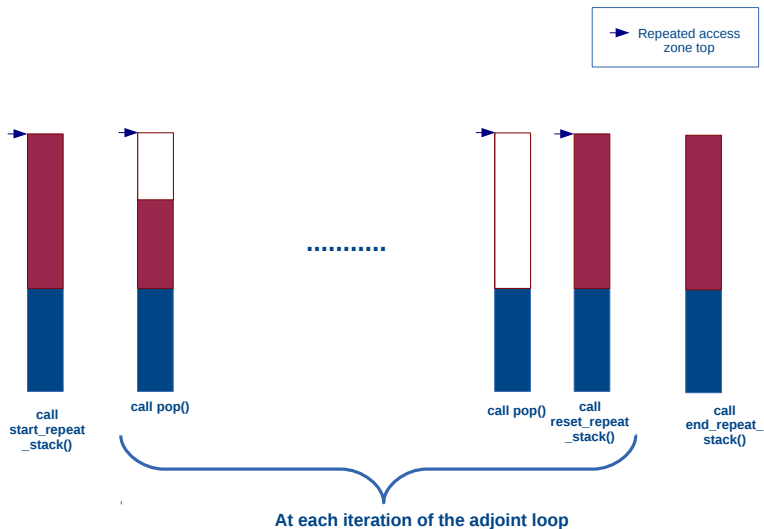
CALL reset_repeat_stack()

end


CALL end_repeat_stack()

$$\bar{x} = \bar{w} * \frac{\partial \phi}{\partial x}(z, x) + \bar{x}$$

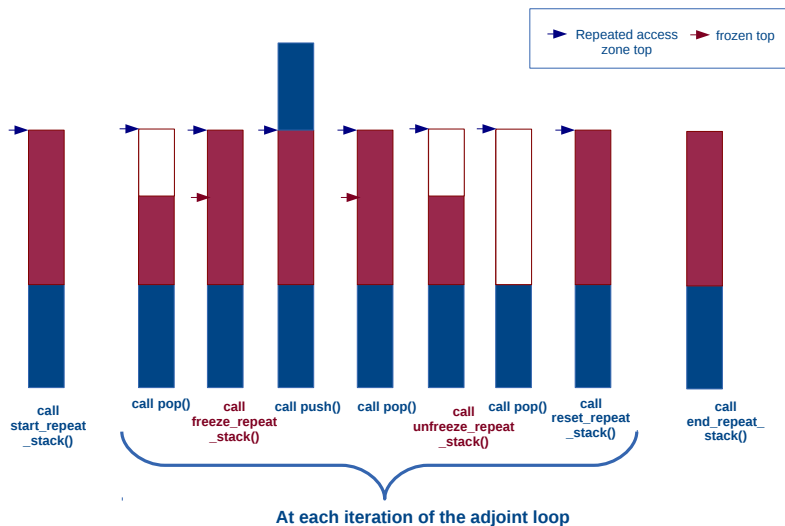
Extension of the stack mechanism



Extension of the stack mechanism



Extension of the stack mechanism




Experiment on real-medium size code

We validated our implementation on a medium-size code called GPDE:

- A Fortran90 program developed at Queen Mary University Of London (QMUL)
- An unstructured pressure-based steady-state Navier-Stokes solver with finite volume spatial discretization
- Contains a Fixed-Point loop that computes the pressure and velocity (the state variables) of an incompressible flow.

Results

- Small improvement of the accuracy of the result



- Major reduction of the memory needed to store the trajectories.

	Two-Phases adjoint	Black Box adjoint
Memory Storage	10.1 MBytes	605.5 MBytes

Nested FP loops: constant initial guess for the inner loop

```

Do while z not converged
  x = g(z, F)
  zin = const
  Do while zin not converged
    zin = Jacobi(zin, M, x)
  end
  z = zin
end
y = f(z)

```

	Two-Phases adjoint	Black Box adjoint
Number of iterations	10132	14491
Accuracy	$12 * 10^{-5}\%$	$2.1 * 10^{-5}\%$
Accuracy: 14491 iterations	$1.1 * 10^{-5}\%$	$2.1 * 10^{-5}\%$
Memory storage	268 bytes	86 Kbytes

Nested FP loops: smart initial guess for the inner loop

Do while z not converged

$$x = g(z, F)$$

Do while z_{in} not converged

$$z_{in} = \text{Jacobi}(z_{in}, M, x)$$

end

$$z = z_{in}$$

end

$$y = f(z)$$

	Two-Phases adjoint	Black Box adjoint
Number of iterations	10132	8788
Accuracy	$12 * 10^{-5}\%$	$2 * 10^{-5}\%$
Memory storage	268 bytes	86 Kbytes

Black Box on a nested FP loops with constant initial guess

Do while z not converged

$x = g(z, F)$

$z_{in} = \text{const}$

Do while z_{in} not converged

$\text{push}(z_{in})$

$z_{in} = \text{Jacobi}(z_{in}, M, x)$

end

$\text{push}(z)$

$z = z_{in}$

end

...

...

Do "as many times"

$\text{pop}(z)$

$\bar{z}_{in} = \bar{z} + \bar{z}_{in} ; \bar{z} = 0$

Do "as many times"

$\bar{z}_{in+} = \bar{z}_{in} * \frac{\partial \text{Jacobi}}{\partial z_{in}}(z_{in}, M, x)$

$\bar{x}+ = \bar{z}_{in} * \frac{\partial \text{Jacobi}}{\partial x}(z_{in}, M, x)$

$\text{pop}(z_{in})$

end

$\bar{z}_{in} = 0$

$\bar{z}+ = \bar{x} * \frac{\partial g}{\partial z}(z, F)$

$\bar{F}+ = \bar{x} * \frac{\partial g}{\partial F}(z, F)$

end

Black Box on a nested FP loops with smart intial guess

Do while z not converged

$$x = g(z, F)$$

Do while z_{in} not converged

push(z_{in})

$$z_{in} = \text{Jacobi}(z_{in}, M, x)$$

end

push(z)

$$z = z_{in}$$

end

...

...

Do "as many times"

pop(z)

$$\bar{z}_{in} = \bar{z} + \bar{z}_{in} ; \bar{z} = 0$$

Do "as many times"

$$\bar{z}_{in+} = \bar{z}_{in} * \frac{\partial \text{Jacobi}}{\partial z_{in}}(z_{in}, M, x)$$

$$\bar{x}+ = \bar{z}_{in} * \frac{\partial \text{Jacobi}}{\partial x}(z_{in}, M, x)$$

pop(z_{in})

end

$$\bar{z}+ = \bar{x} * \frac{\partial g}{\partial z}(z, F)$$

$$\bar{F}+ = \bar{x} * \frac{\partial g}{\partial F}(z, F)$$

end

Nested FP loops: smart initial guess for the inner loop

Do while z not converged

$$x = g(z, F)$$

Do while z_{in} not converged

$$z_{in} = \text{Jacobi}(z_{in}, M, x)$$

end

$$z = z_{in}$$

end

$$y = f(z)$$

	Two-Phases adjoint	Black Box adjoint
Number of iterations	10132	8788
Accuracy	$12 * 10^{-5}\%$	$2 * 10^{-5}\%$
Memory storage	268 bytes	86 Kbytes

Two-Phases on a nested FP loops with smart initial guess

Do while z not converged

$$\mathbf{x} = \mathbf{g}(\mathbf{z}, \mathbf{F})$$

Do while z_{in} not converged

$$\mathbf{z}_{in} = \text{Jacobi}(\mathbf{z}_{in}, \mathbf{M}, \mathbf{x})$$

end

$$\mathbf{z} = \mathbf{z}_{in}$$

end

last iteration + store(\mathbf{z} , \mathbf{z}_{in})

Do while $\bar{\mathbf{w}}$ not converged

restore(\mathbf{z})

$$\bar{\mathbf{z}}_{in} = \bar{\mathbf{w}} + \bar{\mathbf{z}}_{in}; \bar{\mathbf{w}} = 0$$

$$\bar{\mathbf{w}}_{in} = \bar{\mathbf{z}}_{in}$$

Do $\bar{\mathbf{w}}_{in}$ not converged

$$\bar{\mathbf{w}}_{in+} = \bar{\mathbf{w}}_{in} * \frac{\partial \text{Jacobi}}{\partial \mathbf{z}_{in}}(\mathbf{z}_{in}, \mathbf{M}, \mathbf{x})$$

$$\bar{\mathbf{w}}_{in} = \bar{\mathbf{w}}_{in} + \bar{\mathbf{z}}_{in}; \text{restore}(\mathbf{z}_{in})$$

end

$$\bar{\mathbf{x}}+ = \bar{\mathbf{w}}_{in} * \frac{\partial \text{Jacobi}}{\partial \mathbf{x}}(\mathbf{z}_{in}, \mathbf{M}, \mathbf{x})$$

$$\bar{\mathbf{w}}+ = \bar{\mathbf{x}} * \frac{\partial \mathbf{g}}{\partial \mathbf{z}}(\mathbf{z}, \mathbf{F})$$

$$\bar{\mathbf{w}} = \bar{\mathbf{w}} + \bar{\mathbf{z}}$$

end

$$\bar{\mathbf{F}}+ = \bar{\mathbf{x}} * \frac{\partial \mathbf{g}}{\partial \mathbf{F}}(\mathbf{z}, \mathbf{F})$$

Two-Phases with smart initial guess for the inner adjoint

Do while z not converged

$$x = g(z, F)$$

Do while z_{in} not converged

$$z_{in} = \text{Jacobi}(z_{in}, M, x)$$

end

$$z = z_{in}$$

end

last iteration + store(z, z_{in})

Do while \bar{w} not converged

restore(z)

$$\bar{z}_{in} = \bar{w} + \bar{z}_{in}; \bar{w} = 0$$

$$\bar{w}_{in} = \bar{z}_{in} + \bar{w}_{inConv}$$

Do \bar{w}_{in} not converged

$$\bar{w}_{in+} = \bar{w}_{in} * \frac{\partial \text{Jacobi}}{\partial z_{in}}(z_{in}, M, x)$$

$$\bar{w}_{in} = \bar{w}_{in} + \bar{z}_{in}; \text{restore}(z_{in})$$

end

$$\bar{x}+ = \bar{w}_{in} * \frac{\partial \text{Jacobi}}{\partial x}(z_{in}, M, x)$$

$$\bar{w}_{inConv} = \bar{w}_{in} * \frac{\partial \text{Jacobi}}{\partial z_{in}}(z_{in}, M, x)$$

$$\bar{w}+ = \bar{x} * \frac{\partial g}{\partial z}(z, F); \bar{w} = \bar{w} + \bar{z}$$

end

$$\bar{F}+ = \bar{x} * \frac{\partial g}{\partial F}(z, F)$$

Nested FP loops: smart initial guess for the inner loop

```

Do while z not converged
  x = g(z, F)

  Do while zin not converged
    zin = Jacobi(zin, M, x)
  end
  z = zin
end
y = f(z)

```

	Two-Phases with smart initial	Two-Phases	Black Box
Number of iterations	5219	10132	8788
Accuracy	$10 * 10^{-5}\%$	$12 * 10^{-5}\%$	$2 * 10^{-5}\%$
Accuracy for 8788 iterations	$0.98 * 10^{-5}\%$	$340 * 10^{-5}\%$	$2 * 10^{-5}\%$

Summary

- We are seeking to improve performance of adjoint codes produced by the adjoint mode of Automatic Differentiation
- We explained why we consider the strategy initially proposed by Christianson as the best suited for our needs.
- We described the way we implemented this strategy in Tapenade.
- We experimented this strategy on a some a real medium size code and quantified its benefits.
- We studied the related question of the initial guess in the case of nested FP iterations.

Further Work

- The restrictions we imposed on the structure of Fixed-Point loops can be lifted
- Study the choice of the initial guess of the adjoint loop in the general case of nested structures of FP loops.
- Derive the stopping criterion of the adjoint loop from the original loop's stopping criterion

Acknowledgements

This work has been conducted within the **About Flow** project on “Adjoint-based optimisation of industrial and unsteady flows”.

<http://aboutflow.sems.qmul.ac.uk>

About Flow has received funding from the European Union’s Seventh Framework Programme for research, technological development and demonstration under Grant Agreement No. 317006.



National
Technical
University of
Athens