



Published on *CD-adapco* (<http://www.cd-adapco.com>)

[Home](#) > [Unlocking the Mystery of the Adjoint Solver](#)

Unlocking the Mystery of the Adjoint Solver



Sabine Goodwin

June 17, 2013, 5:38 pm

[Share This Post](#)

The biggest challenge of gradient-based shape optimization using high-fidelity CFD has always been the formidable computational expense tied to constructing the sensitivity of the objective (or cost) functions with respect to the design variables.

Take for example the case of a typical external aerodynamics application. An aircraft manufacturer wants to improve on the operational performance of the existing fleet by adding winglets to their aging aircraft. This modernization will only pay off if the installation of the winglet results in significant fuel savings. The company aerodynamics expert is convinced that a good business case can be made and she takes the initiative by designing an initial winglet that can be retrofitted on the aging aircraft. The design is parameterized by a significant number of design variables such as span-wise airfoil sections with varying thickness and camber, winglet sweep, taper ratio, toe angle, cant angle and height. An initial CFD solution at cruise C_L confirms that the aerodynamicist knows her craft because the winglet already shows potential on the first try! Results indicate a slight reduction in overall drag of the system, but unfortunately the wing root bending moment has increased which will require re-enforcing of the structure, negating the drag reduction in the process. The question for our expert now becomes: how does altering the shape of the winglet affect span-wise lift distribution (or induced drag) and how can I design an efficient winglet requiring minimal structural modifications and still meet my tight project deadline?

The standard approach to solving this problem has been to use finite-differences to estimate the gradient of the cost function with respect to the design variables. This is a costly procedure that involves first computing a baseline solution, then perturbing each design parameter one at a time and performing an additional solution following each perturbation; a process that amounts to one additional flow solution for every design variable. A gradient-based optimizer can then use this information, together with repeated evaluations of the cost, to alter the geometry in such a way that the cost function is driven to a minimum. You can imagine that in the case of the winglet design, with such a significant number of design variables defining the shape, the high computational cost associated with this approach makes it almost intractable. In addition, it's also an inefficient way of doing things because the influence of the parameter changes are only understood after multiple iterations of the optimization cycle. The only option for the designer is to significantly decrease the CPU requirements by reducing the number of design variables to just a few, greatly limiting the

scope and success of the project.

What if there was a powerful tool available to compute the sensitivity of the cost function with respect to many design variables at the CPU cost equivalent to just one flow solution? And not only that? What if this method also provided guidance on how to best optimize the design from the start? Integration of an adjoint solver as part of a CFD suite enables this economical sensitivity analysis and is a promising strategy for performing shape optimization involving few cost functions and many design variables, as is the case for the aircraft manufacturer's winglet design for derivative aircraft.

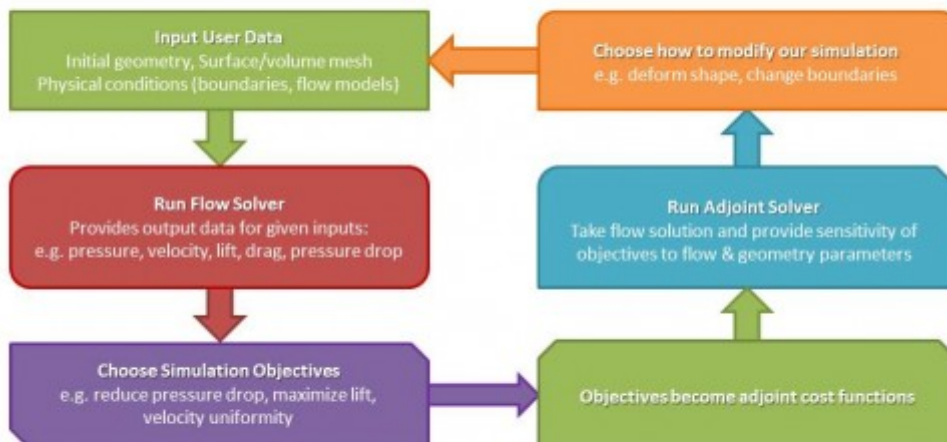
For most of us, adjoint solvers remain a bit of a mystery and the mathematical underpinnings of the method can be quite intimidating. To explain it in very general terms, the approach is based on optimal control theory where the cost function is defined using Lagrange multipliers to include both the flow solution and the mesh movement solution as constraints on the optimization. By taking the derivative of this cost function with respect to the design variables first, and by solving the adjoint equations for the Lagrange multipliers next, the sensitivity of cost function with respect to the flow residual (mass, momentum and energy) and the mesh coordinates is obtained. It is then straightforward to compute the sensitivity of the cost function with respect to all of the design variables simultaneously through a simple matrix multiplication.

The adjoint solver is a very attractive approach because not only does it provide great insight into the system performance early on, it also offers a faster route to an improved design. Let's go back to the winglet design problem. It does not matter how many design variables were required to parameterize the geometry, by performing one solution of the flowfield on the baseline winglet and one solution of the adjoint equations (at a CPU cost equivalent to one flow solution), all the important information is now available for the optimizer to make an informed decision towards improvement. Imagine the cost and time savings!

And it doesn't end here. With the adjoint solver, the sensitivity of the cost function with respect to the flow residual is computed. This is effectively a measure of error in the solution and opens the door for uncertainty quantification. In addition, it provides a mathematical formulation for identifying areas requiring further mesh resolution to better capture the cost function. Instead of depending on solution gradients to resolve the mesh, the aerodynamicist can let the code do the work and she is likely to be impressed by the non-intuitive ways the mathematics refines the winglet mesh. This has the potential of significantly reducing mesh sizes while increasing accuracy resulting in a much more practical optimization process.

Optimal control theory has been around for many years, so why is an adjoint approach to design optimization using CFD not yet main-stream? The bottom line is that implementation with CFD can be complex (there are several approaches e.g. continuous vs. discrete) and especially the transformation of surface sensitivities into a smooth shape through mesh morphing is challenging. The demand from our user community has been high and as a result, STAR-CCM+ will include an integrated discrete adjoint solver in v8.04. The solver provides both 1st and 2nd order adjoints for a wide range of cost functions and is broadly compatible with the existing STAR-CCM+ physics models. CD-adapco has a dedicated team working on delivering additional features for the adjoint solver in the near future including adding the mesh morpher and integration with optimization tools. Look out for it in future releases!

Adjoint Approach to Optimization



Industries:

Aerospace & Defense [1]

Products:

STAR-CCM+® [2]

CD-adapco is the world's largest independent CFD focused provider of engineering simulation software, support and services. We have over 30 years of experience in delivering industrial strength engineering simulation.

Source URL: <http://www.cd-adapco.com/blog/sabine-goodwin/unlocking-mystery-adjoint-solver?language=en>

Links:

[1] <http://www.cd-adapco.com/industries/aerospace-defense>

[2] <http://www.cd-adapco.com/products/star-ccm%C2%AE>