With adjoint equations from simulation to optimisation

Rene Schneider

Abstract:
In recent years the focus is shifting from the simulation of complex systems to their optimisation. This extension is natural, as the desire to simulate and understand a system originates itself from the hope to influence the behaviour of the system.

We discuss a simple but general way of extending the simulation software to an optimisation tool. This is done with simulations based on PDE or ODE solvers in mind, where the solution process typically involves a large number of variables and equations, and often a large number of parameters are available for optimisation as well.

The approach can be summarised as using standard gradient based optimisation algorithms, taking special care of the efficiency of gradient evaluations. These algorithms are known for their fast convergence (in terms of iterations), but the cost of evaluating the gradients may render them inefficient for large problems, if approached the wrong way.

We demonstrate that the cost for the evaluation of a gradient is comparable to only one additional simulation, independent of the number of optimisation parameters. This efficiency is achieved by the use of the discrete adjoint technique, a systematic method for differentiating the simulation code. The technique is closely related to AD (algorithmic or automatic differentiation), which itself can help to allow easy implementation of the approach.

The generality of the approach makes it suitable for very different optimisation problems. As an example of generalised utilisation, we also demonstrate how it can be used in adaptive mesh design for elliptic PDE problems, improving the efficiency of the simulation itself.

Various numerical examples of the optimisation technique will be shown.