Institution: University of Oxford

Unit of Assessment: 10

Title of case study: Adjoint sensitivities in computational finance bring orders-of-magnitude runtime improvements

1. Summary of the impact

The largest investment banks in London each have thousands of servers largely devoted to Monte Carlo simulations, and to quantify their risks and satisfy regulatory demands they need to be able to calculate huge numbers of sensitivities (defined below) known collectively as "Greeks". An adjoint technique developed by Professor Mike Giles in 2006 greatly reduced the computational complexity of these calculations. The technique is used extensively by Credit Suisse and other major banks, reducing their computing costs and energy consumption. It has also led to the Numerical Algorithms Group developing new software to support the banks in exploiting this new adjoint approach to computing sensitivities.

2. Underpinning research

Adjoint techniques are a well-established set of mathematical methods that have been extensively used in engineering design optimisation to simultaneously and efficiently compute the sensitivity of a single output quantity with respect to a large number of input parameters.

Mike Giles (Oxford faculty, 1992 to date) has been a leading researcher in the use of adjoints in engineering design optimisation; his introductory article [1] on the subject with Niles Pierce (Oxford postdoctoral researcher, 1997-1998) has been well cited. When he switched research fields from computational fluid dynamics to computational finance, Giles recognised the opportunity to apply the adjoint technique to Monte Carlo option pricing in finance to more efficiently compute option price sensitivities (known in the industry as "Greeks") which are used to estimate, and thereby minimise, possible future losses due to changes in stock prices, interest rates, exchange rates, etc. In January 2006, together with Professor Paul Glasserman from Columbia University, he published the paper "Smoking adjoints: fast Monte Carlo Greeks" [2] in RISK, the leading publication for those working in quantitative finance within investment banks and other financial institutions. This is the key paper underpinning this Impact case study.

Manual implementation of discrete adjoint methods can be time-consuming and error-prone. Fortunately, much of this can be automated using forward and reverse mode automatic differentiation methods developed in computer science. In this context, one piece of research by Mike Giles was on a higher-level linear algebra approach that is relevant to key steps in Monte Carlo simulation (such as the Cholesky decomposition of a correlation matrix) and time-marching in financial PDE simulations [3].

A key technical limitation was the fact that the whole approach requires differentiability, but many financial option payoffs are discontinuous. This was addressed by inventing the "vibrato" Monte Carlo method [4], which is a hybrid mix of the pathwise sensitivity method (which the adjoint treatment is based on) and the alternative, less efficient, Likelihood Ratio Method.

The specific requirements of correlation Greeks (which are the sensitivity of an option price to two or more inputs simultaneously varied) to compute the sensitivity to each of the many elements in the correlation matrix was addressed in [5], a *RISK* paper written with Luca Capriotti, a Director at Credit Suisse Group, Investment Banking Division, and probably the leading proponent of adjoint techniques within the industry. Forward and reverse mode automatic differentiation was introduced in [6], a further *RISK* paper with Luca Capriotti.

3. References to the research

*[1] MB Giles, NA Pierce. 'An introduction to the adjoint approach to design', Flow, Turbulence and Combustion, 65(3): 393-415, 2000, http://people.maths.ox.ac.uk/gilesm/files/ftc00.pdf. 130 citations according to Web of Science, 280 according to Google Scholar

*[2] MB Giles, P Glasserman. 'Smoking adjoints: fast Monte Carlo Greeks', RISK, 19(1):88-92, January 2006, http://www2.gsb.columbia.edu/faculty/pglasserman/Other/RiskJan2006.pdf.

*[3] MB Giles. 'Collected matrix derivative results for forward and reverse mode algorithmic differentiation'. pp. 35-44 in Advances in Automatic Differentiation, Springer, 2008, http://people.maths.ox.ac.uk/gilesm/files/AD2008.pdf.

[4] MB Giles, 'Vibrato Monte Carlo sensitivities', pp. 369-392 in *Monte Carlo and Quasi Monte Carlo Methods 2008*, Springer, 2009, http://people.maths.ox.ac.uk/gilesm/files/mcqmc08.pdf.

[5] L Capriotti, MB Giles. 'Fast correlation Greeks by adjoint algorithmic differentiation', *RISK*, 23(4):77-83, 2010, http://people.maths.ox.ac.uk/gilesm/files/risk_cg10.pdf.

[6] L Capriotti, MB Giles. 'Algorithmic differentiation: adjoint Greeks made easy', *RISK*, 25(10), 2012, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1801522.

The three asterisked outputs best indicate the quality of the underpinning research. Papers [5] and [6] are refereed papers in one of the leading practitioner publications.

Research funding

Research was partially funded by a 15-month EPSRC Springboard Fellowship: Development of Multilevel Monte Carlo Algorithms for Mathematical Finance, EP/E031455/1, 01/01/07-31/03/08, £74,574.

4. Details of the impact

The research has given rise to significant economic and environmental impact, benefiting a significant percentage of the major (so-called "Tier 1") banks, the software industry that supports the banks, and reducing significantly the energy consumption per Monte Carlo simulation performed by financial institutions.

From research to impact

The largest investment banks in London each have thousands of servers largely devoted to Monte Carlo simulations, and to quantify their risks (the amount of money they could potentially lose due to uncertain and unpredictable events) and satisfy regulatory demands they need to be able to calculate huge numbers of sensitivities efficiently. The "Smoking adjoints" paper [2] was recognised immediately as a significant advance on the state-of-the-art. In January 2007 the paper was voted by the finance industry readers of *RISK* as the best paper of 2006, with Mike Giles and Paul Glasserman being jointly named "Quant of the Year" by the magazine [A]. Quant of the Year is "the award practical quants most care about", according to Laurie Carver of *RISK* magazine who leads the survey of "authors, referees and other industry and academic quants" to decide the winner each year. This indicates the broad importance of this work to the finance industry, as does the fact that Giles has been asked to give a one-day course on "Adjoint Methods for Option Pricing" at Global Derivatives Trading and Risk Management 2013, the leading international conference for the industry.

Nature and extent of the impact

One bank, Credit Suisse, was very quick to adopt the approach and published two papers on the subject. One in RISK in 2011 entitled 'Real Time Counterparty Credit Risk Management in Monte Carlo' [B] builds on the work of [2-6] and states "*Adjoint algorithmic differentiation can be used to implement efficiently the calculation of counterparty credit risk. We demonstrate how this powerful technique can be used to reduce the computational cost by hundreds of times, thus opening the* *way to real time risk management in Monte Carlo*." In 2008, Credit Suisse filed a US patent application (pending) [C] "*to protect our use of it*" and informed us that the adjoint technique is "*very well appreciated within the company because of its benefits to better risk management and enhanced profitability through improved hedging*" and a Director at Credit Suisse Group states that [D] the impact is far reaching as "*it's clear that many banks are now using the methodology*". Support for the broad impact of the "Smoking adjoints" paper [2] comes from the Numerical Algorithms Group (NAG, http://www.nag.co.uk/), a company which develops mathematical libraries and other software. The Vice-President of Sales at NAG talks extensively to the banks. His supporting letter [E] states that "*possibly 20% of the Tier 1 banks are using adjoints on a daily basis as part of their routine option pricing processes, and most of the others have tested and are developing the technology*". He further comments [E] that leading quants have stated that "*because of regulatory and internal risk management concerns, many more Greeks are being required for a wide range of other pricing applications as well, and this is driving the growth in the use of adjoints, to minimise the computational costs of generating all of these*". NAG currently provide consultancy services to banks to help them generate adjoint codes, and in response to requests from banks have also developed adjoint versions of some of their key mathematical library routines using their DCO software which was released in 2010. The Vice-President of Sales at NAG states [E] "*We see this as a growth area for us and we plan further investment next year*."

Ingo Schneider, the head of financial engineering at Dekabank, wrote a paper in *RISK* in 2009 with two academic collaborators on "Fast Monte Carlo Bermudan Greeks" [F]. The abstract of the paper describes it as extending "*the pioneering work of Giles & Glasserman (2006)*". The method was described in *RISK*'s regular technical section [G] and commented upon by an anonymous credit derivatives trader at a European bank: "*It means you can do far more simulations, for far more risk parameters, with far less computational cost. If you want to do it in real time – and with the speed news travels now there's really no other option – there's not really another way to do it*."

Credit Suisse confirm that when used in credit valuation adjustment (CVA), for example, the saving in computing costs is considerable. "*We're talking orders of magnitude … Something that would have taken hours or overnight by bumping can be calculated hundreds of times quicker, and you can manage the CVA intra-day, in real time*." [H].

As indicated by Credit Suisse, the computational benefits of the adjoint approach can be very substantial. For Correlation Greeks, and Greeks for Fixed Income products which depend on a large set of future interest rates, the savings can be up to a factor of 100 compared with the old approach (known as 'bumping') of simply perturbing each input parameter individually and re-simulating to find the consequential change in the value. In practical terms, this means that banks can either reduce their computing costs and energy consumption, or increase the number of simulations they can afford to perform.

The scale of the calculations by the banking sector is indicated by the table below which highlights 8 of the world's top 500 supercomputers [I] that are in the UK and are likely to be performing financial Monte Carlo simulations.

rank	company	size
#185	IT service provider	17,280 cores
#275	financial institution	24,672 cores
#317	IT service provider	10,288 cores
#354	IT service provider	14,592 cores
#417	IT service provider	16,288 cores
#461	financial institution	18,080 cores
#465	bank	17,952 cores
#476	IT service provider	6,560 cores

The IT providers in the list are typically providing computing resources to large financial institutions.

Note also that the computing facilities of many banks are probably not included in this list. The information is provided by the IT vendors (banks are very secretive about their computing facilities).

An indication of the power consumption of financial data centres is given by a 2008 Guardian article (http://www.guardian.co.uk/technology/2008/may/29/energy.olympics2012) on the negative impact that the Olympics would have on power provision for the banks based in Canary Wharf. What happened subsequently is that major banks built new data centres around the M25 where power was available, or used new data centres set up by IT providers. This indicates the computational power of the major financial institutions, and it is thought that over half of this power is devoted to Monte Carlo simulations. Given that the costs of renting supercomputer time is upwards of thousands of dollars per hour [J], in this context, even a factor 2 reduction in the cost of computing Monte Carlo sensitivities equates to a major reduction in cost as well as energy consumption, or enables many more calculations giving improved modelling and mitigation of risk through considering many more different risk scenarios.

5. Sources to corroborate the impact

[A] Risk Quant of the Year award: http://www.risk.net/risk-magazine/feature/1498251/quants-paul-glasserman-michael-giles

[B] L Capriotti, J. Lee, M. Peacock, 'Real Time Counterparty Credit Risk Management in Monte Carlo', RISK 24(6):86-90, 2011 http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1824864&rec=1&srcabs=1801522

[C] Credit Suisse US patent application, based on Giles' adjoint techniques: https://www.google.com/patents/WO2008151098A1?cl=en&dq=luca+capriotti&hl=en&sa=X&ei=9O4EUsSFEuqG0AWq24FI&sqi=2&pjf=1&ved=0CDsQ6AEwAQ.

[D] Letter from a Director at Credit Suisse Group, Investment Banking Division, describing the nature and extent of the impact at Credit Suisse. Copy held by the University of Oxford, 2013

[E] Letter from the Vice-President, Sales, NAG, describing the nature and extent of the impact on NAG and the financial sector. Copy held by the University of Oxford, 2013

[F] M. Leclerc, Q. Liang, I. Schneider, 'Fast Monte Carlo Bermudan Greeks', RISK, 22(7):84-88, 2009, http://people.maths.ox.ac.uk/gilesm/files/risk_greeks09.pdf

[G] L Carver Cutting Edge introduction: Computation, computation, computation', RISK Technical paper, 06 Sep 2012, commenting on the significance of Adjoint methods in the financial industry: http://www.risk.net/risk-magazine/technical-paper/2203043/cutting-edge-introduction-computation-computation-computation

[H] L Carver, 'Algorithmic gymnastics', RISK, 25(8):52, 2012, http://www.risk.net/risk-magazine/profile/2194286/credit-suisse-algorithmic-gymnastics

[I] The Top 500 supercomputer list http://top500.org/ (downloaded on 3/9/13), showing extent of computational resources used by banks

[J] Price of supercomputing time can be verified at: http://arstechnica.com/business/2012/04/4829-per-hour-supercomputer-built-on-amazon-cloud-to-fuel-cancer-research/

References [B], [F] and [H] are by practitioners, indicating the reach of Giles' adjoint research