Institution: University of Hertfordshire		
Unit of Assessment: 11 – Computer Science and Informatics		
Title of case study: Automatic Differentiation: faster, more accurate modelling to strengthen risk management in financial services and optimise engineering performance		
Period when the underpinning research was undertaken: 2000 - 2013		
Details of staff conducting the underpinning research from the submitting unit:		
Name(s):	**Role(s) (e.g. job title):**	**Period(s) employed:**
Bruce Christianson	Professor	1987 – 2018
Uwe Naumann	Senior Lecturer then Visiting Research Fellow	2000 – 2001 2001 – present
Jan Riehme	Research Assistant then Senior Research Fellow	2002 – 2003 2009 – 2011
Period when the claimed impact occurred: 2014 – 31 Dec 2020		
Is this case study continued from a case study submitted in 2014? Y		

1. Summary of the impact (indicative maximum 100 words)

A decade of EPSRC-funded research led by the University of Hertfordshire (UH) underpinned the development of novel algorithmic methods that significantly accelerated and improved the accuracy of risk calculations for investment banks and design sensitivities in engineering. The full commercialisation of these Automatic Differentiation (AD) tools in 2014 by Oxford-based Numerical Algorithms Group (NAG) Ltd has:

- Strengthened risk management in derivatives trading for [TEXT REDACTED FOR PUBLICATION], enabling more informed investment decisions and more efficient regulatory adherence.
- Resulted in two leading banks winning Risk Awards for implementation of the tools; one credited the algorithmic techniques for the avoidance of hundreds of millions of pounds in losses during the Covid-19 pandemic and an increase in trade flows that contributed to a 33% revenue rise.
- Optimised aerodynamic performance for a Formula One team and a solar car that completed the World Solar Challenge 2019.
- Added new capabilities to a tool used to optimise the design of marine infrastructure.
- [TEXT REDACTED FOR PUBLICATION]

2. Underpinning research (indicative maximum 500 words)

Society is increasingly reliant on the numerical simulation of complex real-world phenomena. In many industries it has become important not only to compute the simulation value, but also its sensitivity to the model inputs in a highly accurate and computationally efficient way. Automatic (also known as Algorithmic) Differentiation (AD) is a set of mathematical techniques for transforming numerical modelling code so that it calculates the numerical sensitivities of the model values as well as the values themselves, and to the same precision. The adjoint (or reverse) mode of AD (AAD) involves making the program run backwards and recovering all the intermediate numerical values in the process; for some models this mode can compute sensitivities thousands of times faster than finite difference methods. In finance for example, the use of AAD can significantly speed up the calculation of sensitivities of derivatives prices to underlying factors, known as Greeks, i.e. the different dimensions of risk involved in taking an options position.

Research at UH's Centre for Computer Science and Informatics Research (CCSIR) in the early 2000s made a significant contribution to advancing the theory underpinning the technology. Led by Christianson, the UH team expanded the basic notions of AD and identified lines of research that could lead to the development of practical tools for optimisation [**3.1**]. Christianson then worked with engineering researchers at Cranfield University to apply AD to race car performance optimisation. They showed that use of an AD software tool to calculate the optimal vehicle controls (i.e. steer angle and driving/braking torque) for a simple, one-turn manoeuvre was more robust and ten times faster than using finite difference methods [**3.2**]. In 2002, CCSIR began a 12-year collaboration with RWTH Aachen University and Oxford-based Numerical Algorithms

Group (NAG), a numerical software and HPC services provider. Funded by four EPSRC grants, the 'CompAD' project, led by Christianson, succeeded in developing the world's first industrial-strength Fortran compiler with built-in support for AAD.

Under the first grant [**G1**], the research group sought to ease the burden that AD placed upon the user (at the time, AD tools required intensive, time-consuming configuration) by providing new datatypes and creating easy-to-use interface routines for extracting the derivatives. Having had the idea to place AD functionality inside a Fortran compiler [**G2**], the group succeeded in integrating AD capabilities into the NAGWare Fortran 95 Compiler and then, for the first time, embedded the ability to generate adjoint code automatically [**3.3, 3.4**]. This was adequate for small to medium-sized problems (up to a few hundred input variables). However, the solving of even moderately large problems (thousands of input variables) required the systematic use and manual placement of checkpoints in order to manage the trade-off between storage on the way forward and recomputation on the way back. This was considered limiting and time-consuming for users without previous experience of AD and represented a barrier to the uptake of numerical methods based upon it. Therefore, the third project [**G3**] automated this process of trading off storage and recomputation in a way that was close to optimal. The AD-enabled NAG Fortran Compiler was used to generate a discrete adjoint version of a geophysical fluid dynamics numerical model [**3.5**]. Higher order derivatives were generated automatically by feeding back into the compiler parts of its own output during the compilation process [**3.6**].

The final project [**G4**] sought to turn the tool into commercially viable AD software. The team achieved this by extracting the AD functionality from the CompAD compiler and integrating it with the NAGWare Fortran Library, a set of algorithms designed to solve complex mathematical problems quickly and easily. Proof of its efficiency and viability was demonstrated through an application for the German Waterways Board to improve dredging along the River Danube. CompAD was used to predict the flow and profile of sediment. Instead of calculating just a single value at one point in time and space, it allowed a probability distribution of the evolution to be automatically derived. A reliability analysis then ranked uncertainty input parameters.

3. References to the research (indicative maximum of six references)

3.1 Bartholomew-Biggs M, Brown S, Christianson B, Dixon L. Automatic Differentiation of Algorithms. Journal of Computational and Applied Mathematics. 2000;124(1-2):171-190. https://doi.org/10.1016/S0377-0427(00)00422-2
3.2 Christianson B, Casanova D, Sharp RS, Final M, Symonds P. Application of automatic differentiation to race car performance optimisation. In Corliss G, Faure C, Griewank A, Naumann U, editors, Automatic Differentiation of Algorithms: From Simulation to Optimization. Springer. 2002. p. 117-124 https://doi.org/10.1007/978-1-4613-0075-5_12
3.3 Naumann U, Maier M, Riehme J, Christianson B. Automatic first- and second-order adjoints for truncated Newton. In Ganzha M, Paprzycki M, Pelech-Pilichowski T, editors, Procs of the Int Multiconference on Computer Science and Information Technology (IMCSIT): Vol. 2: Computer Aspects of Numerical Algorithms. 2007. p. 541–555.
3.4 Stumm P, Walther A, Riehme J, Naumann U. Structure-Exploiting Automatic Differentiation of Finite Element Discretizations. In Bischof CH, Bücker HM, Hovland P, Naumann U, Utke J, editors, Advances in Automatic Differentiation. Lecture Notes in Computational Science and Engineering, vol 64. Springer, Berlin, Heidelberg. p. 339-349. https://doi.org/dr23n2
3.5 Rauser F, Riehme J, Leppkes K, Korn P, Naumann U. On the use of discrete adjoints in goal error estimation for shallow water equations. Procedia Computer Science. 2010;1(1):107-115. https://doi.org/10.1016/j.procs.2010.04.013
3.6 Christianson B. A Leibniz notation for automatic differentiation. In Recent Advances in Algorithmic Differentiation. Springer. 2012. p. 1-9. (Lecture Notes in Computational Science and Engineering). https://doi.org/10.1007/978-3-642-30023-3_1

Key underpinning grants

G1 EPSRC: 'Differentiation Enabled Compiler Technology' (CompAD-I). GR/R55252/01. 2002 – 2003. PI: Uwe Naumann. £51,612.
G2 EPSRC: 'Differentiation-Enabled Fortran 95 Compiler Technology' (CompAD-II). EP/D062071/1. 2006 – 2008. PI: Bruce Christianson. £317,665.
G3 EPSRC: 'Differentiation-Enabled Compiler Technology' (CompAD-III). EP/F069383/1. 2008 – 2011. PI: Bruce Christianson. £236,783
G4 EPSRC: 'Towards Industrial Strength Automatic Differentiation' (CompAD-IV). EP/J013358/1. 2012 – 2014. PI: Bruce Christianson. £101,465.

4. Details of the impact (indicative maximum 750 words)

The CompAD compiler, developed through contiguous EPSRC grants, was fully commercialised by NAG in June 2014 with the release of AD-native versions of its NAG optimisation library, the world's largest collection of commercially available numerical and statistical algorithms [**5.1**]. In collaboration with RWTH Aachen and Naumann (who is also a Visiting Research Fellow at UH), NAG built on the CompAD research insights to create dco/c++, an AD-specific software tool for computing sensitivities of C++ codes, and aligned it with its AD-enabled library. NAG further developed its AD/AAD tools over the impact period, with a focus on optimising its offering for the financial services industry. Major updates that made AD/AAD processes quicker, more efficient and more productive were announced at the 14th WBS Quant Finance Conference in September 2018 and in November 2020 [**5.2**].

Direct economic impact via NAG's commercialisation of the CompAD compiler

NAG Ltd is a not-for-profit company limited by guarantee, with offices in Oxford (headquarters), Manchester, Chicago and Tokyo. The CompAD project allowed NAG to offer the benefits of adjoint AD (i.e. greater accuracy, efficiency and robustness in computing sensitivities) to a range of industries, opening up new market opportunities [**5.3**, **5.4**]. Key applications were finance, where rapid computation of sensitivities is used to hedge risk and assess how robust investment portfolios are to market changes; aerodynamics, for example in Formula 1 where adjoint methods allow teams to achieve desired levels of downforce and drag; and oil and gas exploration, where adjoint methods can be used to optimise seismic imaging [**5.4**].

In addition to the licensing of its dco/c++ tool, NAG provides consultancy services to banks to help them use the dco tool to generate adjoint codes. NAG's decision to develop adjoint versions of its library routines was made in response to banks' requests; the company runs workshops to support the quant industry in practical implementation. While non-disclosure agreements limit what can be said, [TEXT REDACTED FOR PUBLICATION] [**5.3**]. [TEXT REDACTED FOR PUBLICATION] [**5.3**].

According to NAG's Chief Commercial Officer, [TEXT REDACTED FOR PUBLICATION] He wrote: *'Our involvement with the CompAD project was a pivotal experience for us in developing a route to exploit the potential of AD technology … We are now broadly recognised as market leader of AD technology. Maintaining a reputation for technical innovation and excellence is not easy especially in financial services where the technology arms race does not stop* [**5.3**]*.'*

Strengthening risk management in financial services for better investment decisions and regulatory adherence – and significant loss avoidance

In financial trading 'Greeks' refer to the dimensions of risk for an options position. They are used by traders and portfolio managers to hedge risk and understand how their profit and loss ledger will behave as prices move. Calculating derivatives exposure across a portfolio is complex. The traditional approach to calculating Greeks has been to make small adjustments to the values of the inputs in the pricing of a derivative and calculate the output value each time, in a process known as bumping. It comes at a significant computational cost, especially for a portfolio of thousands of trades. In many cases, these calculations are prohibitively expensive and cannot be completed within a practical timeframe. Adjoint AD can compute Greeks up to 1,000 times faster than bumping, and with machine accuracy [**5.4**]. Banks use NAG's AD tools to calculate, at far greater speeds and accuracy, sensitivities to hedge risk, explain exposures and assess how quickly portfolios could lose value in a downturn [**5.4**].

Banks tend to be secretive. However, two banks in particular have spoken publicly about their use of AAD and these act as representative case studies for how the wider financial services industry has benefitted. According to a 2019 article in WatersTechnology [**5.5**], Canada's third largest bank Scotiabank (assets: £571bn) used NAG's AD tools to build a new 'risk engine' that is capable of calculating valuation adjustments (XVAs) more accurately and 30 times faster. XVAs take into account the funding, credit and capital costs associated with trading derivatives; traders incorporate XVA into the price of a new derivative trade. Running XVA calculations is one of the largest computational challenges that banks face. The bank said that this new platform '*allows brokers to deliver more accurate derivatives pricing in 20 seconds, which would previously have taken 10 minutes*'. It also '*allows for more nuanced risk analysis thanks to more detailed risk scenario modelling that can assess more than 10 times the number of previous scenarios*' [**5.5**]. As Scotiabank notes, the need to accelerate these calculations has become increasingly important in order to adhere to more stringent international regulations (the Basel III regulatory standards, published December 2017) that are designed to protect the global economy from excessive risk-taking by banks [**5.5**]. AAD has made it possible for banks to calculate XVAs for each new transaction carried out with a client and accurately assess the impact of these adjustments on the bank's balance sheet at the end of every day's trading [**5.5**].

Having worked closely with NAG to develop the AAD algorithm library, Scotiabank 'switched on' the AAD function of its new risk platform in early 2020 [**5.6**]. The result was '*a tool that allows traders and risk managers to better understand exposure across asset classes, including rates, credit, foreign exchange and commodities*' [**5.6**]. This insight proved '*invaluable*' for Scotiabank when the markets plunged in March 2020 amid the Covid-19 pandemic [**5.6**]. The new technology platform enabled the bank to quickly run multiple scenario analyses and implement a hedging strategy that responded to market volatility. The bank's head of XVA trading said: '…*this warned us ahead of time that if we saw a 100 basis point move in credit, we would need to put on large hedges in the rates and FX markets*' [**5.6**]. Major investment banks experienced losses amounting to hundreds of millions during this period (JP Morgan recorded a $951m loss in three months to March 31) [**5.6**]. However, Scotiabank avoided such losses; it did not report an XVA charge on its balance sheet and the bank said the new risk engine had directly increased trade flows, resulting in a 33% year-on-year rise in revenues to the end of October 2020 [**5.6**]. Scotiabank won Technology Innovation of the Year in the prestigious Risk Awards 2021 (which evidences the benefit delivered to the bank during the REF 2021 impact period) [**5.6**].

Danske Bank, Denmark's largest bank, won the In-House System of the Year 2015 Risk award for its implementation of AAD in its production and regulatory systems [**5.7**]. One of the designers of the system authored *Modern Computational Finance: AAD and Parallel Simulations (2018)*, a practical guide for risk professionals on the application of AAD to finance. The book refers to AAD as '*arguably the strongest addition to numerical finance of the past decade*' and now '*a key skill for all quantitative analysts, developers, risk professionals or anyone involved with derivatives*' [**5.7**]. In October 2020, Danske Bank announced it had combined AAD with neural networks to carry out valuation adjustments '*thousands of times faster*' [**5.8**].

Achieving aerodynamic gains and optimising infrastructure design in engineering

Through NAG products, engineering companies have, during the impact period, employed AD to calculate sensitivities of drag, lift or ballistic profile in order to identify the desired geometric shape of a vehicle for optimal performance. NAG reports that a high-profile Formula One team used its adjoint AD tools to improve car performance by designing precise car geometry to achieve desired levels of downforce and drag [**5.3**]. Team Sonnenwagen, a group of students from Aachen University, used NAG's AD-based dco/c++ tool to optimise the design of its solar car that participated in the 2019 World Solar Challenge, a 3022-kilometre race through the Australian outback for vehicles powered only by the sun. The dco/c++ tool was described as '*pivotal*' in identifying design changes that minimised air resistance and '*new suggestions for the cars aerodynamics that would otherwise have gone unnoticed were detected and implemented*' [**5.9**]. Team Sonnenwagen went on to finish sixth out of 53 teams from around the world.

In 2016, the CompAD compiler was incorporated directly into TELEMAC (now TELEMAC-MASCARET), an open-source suite of solvers used to model free-surface and groundwater flow, which is managed by a consortium including French energy company EDF and the German Waterways Board [**5.10**]. This software is used to analyse river flows, sediment transport and wave propagation in the context of marine infrastructure development like dams, bridges and groynes, dredging and flood management. The embedding of AD capabilities in TELEMAC has allowed users of the software to solve previously unsolvable or computationally very expensive problems, thus making a wider range of hydraulic modelling applications possible [**5.10**].

5. Sources to corroborate the impact (indicative maximum of 10 references)

5.1 Announcement at the International Supercomputing Conference of the new AD/AAD tools and services released by NAG, June 13, 2014.
https://news.cision.com/numerical-algorithms-group--nag-/r/new-algorithmic-differentiation-tools-and-services-released-by-nag-and-aachen-collaboration,c9601480
5.2 Announcements of the launch of significant new updates to NAG's AD/AAD products and services in September 2018 and in November 2020.
https://www.nag.com/content/nag-showcases-cva-scale-work-and-new-algorithmic-differentiation-software; https://www.nag.com/news/new-faster-data-fitting-solver-strengthens-nag-library-optimization-suite
5.3 Corroborating statement from the Chief Commercial Officer, NAG Ltd.
5.4 Explanation of NAG's Algorithmic Differentiation Solutions on the NAG website:
https://www.nag.com/content/algorithmic-differentiation-solutions
5.5 '*Scotiabank turns to cloud GPUs for risk calculations'* – article (July 2019) in WatersTechnology citing Scotiabank's use of NAG Ltd's AD software tools.
https://www.waterstechnology.com/data-management/4465306/scotiabank-turns-to-cloud-gpus-for-risk-calculations (full article behind paywall – screenshots on file)
https://www.nag.co.uk/content/impressive-results-scotiabank-using-nag-library-ad-tools-origami-and-azure-xva
5.6 '*Technology innovation of the year: Scotiabank*'. Article on Risk.net detailing why Scotiabank won a Risk Award 2021 for its new risk engine in which NAG's AD algorithms are embedded (while the article was published in February 2021, it corroborates impact realised in 2020):
https://www.risk.net/awards/7736276/technology-innovation-of-the-year-scotiabank
5.7 *Modern Computational Finance: AAD and Parallel Simulations* (2018). Authored by Antoine Savine, a financial derivatives practitioner with Superfly Analytics at *Danske* Bank.
https://www.wiley.com/en-gb/Modern+Computational+Finance:+AAD+and+Parallel+Simulations-p-9781119539452
https://www.wiley.com/en-us/Modern+Computational+Finance:+AAD+and+Parallel+Simulations-p-9781119539452
5.8 '*Danske quants discover speedier way to crunch XVAs*'. Article on Risk.net, October 2020.
https://www.risk.net/our-take/7691456/danske-quants-discover-speedier-way-to-crunch-xvas
5.9 Team Sonnenwagen's use of NAG's AD tools for 2019 World Solar Challenge.
https://www.nag.com/content/team-sonnenwagen-optimize-solar-racing-car-aerodynamics-using-nag-algorithmic-0
5.10 Corroboration of the embedding of Comp-AD compiler in Telemac-Mascaret, March 2016:
http://www.opentelemac.org/index.php/feature-log/217-new-features-of-telemac-2d-version-7-2