

<b>Institution: Imperial College London</b>
<b>Unit of Assessment: 11 Computer Science and Informatics</b>
<b>Title of case study: Case Study 2: Reconfigurable Computing for High Performance Applications (Reconfigurable Computing)</b>
<p><b>1. Summary of the impact</b> (indicative maximum 100 words)</p> <p>In the last 20 years, reconfigurable technology has transformed High-Performance Computing and Embedded Systems Design. Research of the Custom Computing and Reconfigurable Systems groups at Imperial made pivotal contributions to this transformation, targeting particularly Field-Programmable Gate Array (FPGA) technology. Since 2008, the impact of this research has been to</p> <ol style="list-style-type: none"> <li>1. underpin design flow for partial run-time reconfigurable designs for Xilinx FPGA devices;</li> <li>2. contribute to the start-up company Maxeler, pioneering reconfigurable computing systems and cloud services for high-performance computing in the financial and other sectors;</li> <li>3. enable near real-time risk analysis for JP Morgan's global portfolio to analyse and manage risk much faster than previously possible;</li> <li>4. achieve about 250 times speedup for Chevron's seismic modelling for oil and gas exploration, compared to the alternative use of CPU-based machines;</li> <li>5. accelerate a financial market integrity platform for BlueBee and HL Steam in hardware.</li> </ol>
<p><b>2. Underpinning research</b> (indicative maximum 500 words)</p> <p>Field-Programmable Gate Arrays (FPGAs) are a technology introduced in the 80's. They provide an affordable vehicle for hardware acceleration of sophisticated software. The emergence of FPGAs has enabled novel computing engines for solving complex scientific and commercial problems. Research groups in EEE and the Department of Computing (DoC) conducted the underpinning research for the above impacts. The group at EEE, consisting of Prof. Peter Cheung (lead), Dr. George Constantinides and Dr. David Thomas, focused mainly on hardware architecture, synthesis algorithms and applications in digital signal processing. The DoC group, consisting of Prof. Wayne Luk (lead) and Dr. Oskar Mencer, focused mainly on application-specific programming languages, compilers, run-time reconfigurability and system aspects. These research groups and their collaborations have produced significant research advances that underpin the above impacts, clustered in five topics and mapped to references and the above impacts:</p> <p><b>Design of FPGA-based reconfigurable computing systems [1; leading to impact 1]:</b> FPGA-based computing systems provide the flexibility of run-time reconfigurability. However, this advantage also adds a new dimension to the complexity of system design, especially when the FPGA device supports partial run-time reconfiguration. Since 1995, our teams have developed new automatic design tools, which efficiently exploit the dynamic reconfigurable nature of FPGAs in producing novel implementations of computing systems. For specific classes of devices, the techniques we developed reduce reconfiguration time from time linear in their size to constant time at best and logarithmic time at worst.</p> <p><b>An FPGA-specific stream compiler (ASC) [2,4; leading to impacts 2-4]:</b> This stream compiler provides a software-like programming interface to hardware design; it targets FPGAs but maintains the performance of hand-designed circuits. This compiler and interface improve productivity by letting programmers optimize implementations at <i>multiple</i> levels – algorithms, architecture, arithmetic, and gates – and all within the <i>same</i> C++ program. This increased productivity is demonstrated in hardware acceleration of a range of applications, e.g. encryption with Kasumi and IDEA, function evaluation, and Wavelet and LZ-like compression.</p> <p><b>Optimizing word-length and data for area, power, performance, and accuracy trade-offs [3; leading to impacts 2-5]:</b> In conventional computers, the word-length (e.g. in a 32-bit or 64-bit ALU) and data format (e.g. fixed or floating point) are fixed by the architecture. FPGAs provide the freedom of optimally determining word-length and data format according to the needs of specific applications. Based on an approach that combined analytical and heuristic methods, our teams</p>

**Impact case study (REF3b)**

pioneered – since 1999 – a number of techniques for determining word-lengths and data formats that will both optimize area, performance and power consumption for a given system-level specification (e.g. worst-case accuracy and signal-to-noise ratio). These techniques also allow us to trade off the accuracy of computations with other desired performance characteristics (e.g. performance).

**Transferring reconfigurable FPGA-based technology into applications [2-5; leading to impact 3, 4]:** Since 1995, our teams have researched how the aforementioned new techniques can be applied to solve industrial and commercial problems. Our research has demonstrated the applicability of these techniques in various industrial application domains such as financial analysis, seismic modelling and digital signal processing. In particular, we are responsible for some of the earliest research on accelerating financial applications, which produced a powerful mathematical framework for expressing a wide class of Monte-Carlo simulations. In the framework, Monte-Carlo applications can be written in a high-level language, which a streaming compiler such as ours can nonetheless convert into high-performance data pipelines. Our research teams also identified generic design templates that can help reduce design efforts of future systems in these application domains.

**Novel tools for heterogeneous multiprocessor systems [6; leading to impact 5]:** We are among the first (starting this research in 2006) to develop a compilation tool chain for high-level programs targeting heterogeneous systems with different types of processing elements such as general-purpose processors, digital signal processors, and FPGAs. The core of these tools includes a task transformation engine, a mapping selector, an optimizer for data representations, and a hardware synthesizer. The tool chain uses as intermediate representation C programs enriched with source annotations, thus making it easy for users to comprehend and control the compilation process.

**3. References to the research** (indicative maximum of six references)**Publications that directly describe the underpinning research**

\* References that best indicate quality of underpinning research.

\*[1] W. Luk, N. Shirazi and P.Y.K. Cheung. Compilation tools for run-time reconfigurable designs. IEEE Symp. on FPGAs for Custom Comp. Machines (FCCM), pp. 56-65, 1997.

<http://dx.doi.org/10.1109/FPGA.1997.624605>

[2] O. Mencer, D.J. Pearce, L.W. Howes and W. Luk. Design space exploration with a Stream Compiler. IEEE International Conference on Field Programmable Technology (FPT), pp. 270-277, 2003. <http://dx.doi.org/10.1109/FPT.2003.1275757>

\*[3] G. Constantinides, P.Y.K. Cheung and W. Luk. Optimum and heuristic synthesis of multiple word-length architectures. IEEE Trans. on VLSI Systems, 13(1): 39-57, 2005.

<http://dx.doi.org/10.1109/TVLSI.2004.840398>

\*[4] O. Mencer. ASC, A Stream Compiler for Computing with FPGAs. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 25(9): 1603-1617, 2006.

<http://dx.doi.org/10.1109/TCAD.2005.857377>

[5] J.A. Bower, D.B. Thomas, W. Luk and O. Mencer. A reconfigurable simulation framework for financial computation. International Conference on Reconfigurable Computing and FPGAs (ReConFig), pp. 1-9 2006. <http://dx.doi.org/10.1109/RECONF.2006.307750>

[6] W. Luk, J.G.F. Coutinho, T. Todman, Y.M. Lam, W. Osborne, K.W. Susanto, Q. Liu and W.S. Wong. A high-level compilation toolchain for heterogeneous systems. IEEE International System-On-Chip Conference (SOCC), pp. 9-18, 2009. <http://dx.doi.org/10.1109/SOCCON.2009.5398108>

**Grants that directly funded the underpinning research**

[i] HARTES – Holistic Approach to Reconfigurable Real Time Embedded Systems. EU FP6-035143, W. Luk (PI), €764,000, September 2006 – August 2009.

[ii] Optimising Reconfigurable Hardware Parallelism for Concurrent Programs, EPSRC GR/N66599/01, W. Luk (PI), £211,506, April 2001 – December 2004.

## Impact case study (REF3b)

[iii] Effective Production of Parametrised FPGA Libraries, EPSRC GR/L24366/01, W. Luk (PI), £196,706, November 1996 – October 1999.

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

We now provide details of the five aforementioned impacts and their link to underpinning research:

**Impact 1:** Xilinx Corp. is a US\$2.2 billion company (2012) with nearly a 50% share of the FPGA devices market. The impact of our research in FPGA technology is evidenced in their testimony (Distinguished Engineer from Xilinx [A]):

*“The research group led by you (Prof. Luk) in Department of Computing and Professor Peter Cheung in Department of EEE ... is second-to-none internationally, both in terms of quality and quantity. Your team has pioneered many key technologies which have resulted ... in significant industrial impact.”*

For example, the Engineer from Xilinx refers specifically to our work on partial run-time reconfiguration [1], which “... has had a major impact on multiple generations of Xilinx devices that support run-time reconfiguration”, and “..... underpins the design flow for the latest Xilinx device supporting partial run-time reconfiguration”. The specific devices introduced since 2008 were Virtex-5 (65nm), Virtex-6 (40nm) and Virtex-7 (28nm) devices as well as their latest Zynq device family, which embeds ARM processor cores within the FPGA fabric. Xilinx’s document UG702, entitled “Partial Reconfiguration User Guide” and published in October 2012, demonstrates that our work [1], first published in 1997, has stood the test of time: its impact is still found in design flow for Xilinx Zynq devices introduced in 2011 [A].

**Impact 2:** Our underpinning research into acceleration of compute-intensive algorithms using FPGA technology has led to its successful commercial exploitation in the start-up company, Maxeler. The company, with Head Office in Hammersmith, employs 70 members of staff in the UK and US. Many of the key technologies it employs can be traced directly to the research conducted by the team at Imperial College [2, 3, 4]. In particular, the MaxCompiler and MaxGen tools are rooted in Imperial tools. These tools enable MaxCloud, the industry’s first FPGA-based cloud computing service, to provide high performance dataflow computation capability through the cloud [B]. Underpinning research by Cheung and Luk also led to several US patents (US6369610, US7543283, US12/747650) assigned to Maxeler in Nov 2012. Moreover, Maxeler US patent US20130139122 A1 refers to this work on word-length optimisation [3].

**Impact 3:** The FPGA-based technology from Maxeler is used by JP Morgan [C] and by other finance companies such as Scottish Widows [D] for accelerating a variety of financial modelling calculations including risk analysis. This is therefore indirect impact of our research in [2-5]. In one example, the compute time was reduced from 8 hours to less than 4 minutes [E]. The effectiveness of this technology resulted in JP Morgan purchasing a 20% stake in Maxeler [F]. According to JP Morgan’s Managing Director and Global Head of their Applied Analytics Group, the total investment made by JP Morgan in deploying our reconfigurable technology amounts to US\$30m with an annual estimated saving in running cost of US\$6m, in addition to compute time reduction. Furthermore, the reconfigurable computing systems “support our business units that manage the risk on positions in the trillions of dollars” [G].

**Impact 4:** In addition to financial applications, our reconfigurable computing technology [2-5] has had a strong impact on oil and gas exploration. The oil and gas industry is a major user of high-performance computing. In geoscience, computational cycles are dominated by relatively few and well defined kernels. Using Maxeler’s FPGA-based hardware platform and optimising the algorithm implemented with the Maxeler tool flow, speedups of almost 250 times compared to the use of a single CPU core have been reported by Chevron [H]. Faster seismic analysis and imaging allows more application runs or modelling with higher fidelity, or both. Quicker and more accurate results enable oil and gas companies to make more informed bids on parcels and subsequent drilling. Shorter turnaround time on these seismic applications is critical to the company’s profit, especially when they are in a bidding process for drilling rights.

**Impact 5:** BlueBee Technologies is a start-up company that provides tools for heterogeneous multi-core platforms. The success of its tools was recognised by a prestigious Valorisation award from the Dutch Technology Foundation, STW, in December 2012. According to BlueBee’s CEO [I],

## Impact case study (REF3b)

several key components in the BlueBee tool chain – such as task transformation and mapping selection – are directly based on our research on high-level compilation tool chains for heterogeneous systems [6]; success of BlueBee can be attributed to our pioneering research which laid the foundations for many important developments in word-length optimisation [3] and high-level design transformation [6]. Recently BlueBee has teamed up with HL Steam, a company responsible for the Ancoa financial market integrity platform, to support FPGA-based acceleration for such platforms. The impact of our research is therefore also ongoing: the enabling of real-world workloads of millions of financial transactions per second to be efficiently processed at lower dollar and energy cost [J].

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

[A] Distinguished Engineer, Xilinx Research Labs, stating the impact of our research on FPGA technology, particularly on Xilinx products.

[B] “Maxeler launches Maximum Performance Cloud Computing with Maxcloud”, 20 Sept 2011: <http://www.maxeler.com/maxeler-launches-maximum-performance-cloud-computing-with-maxcloud/> Archived on 22/10/2013 <https://www.imperial.ac.uk/ref/webarchive/gyf>

[C] S. Shah, “JP Morgan goes live with Maxeler supercomputer”, 20 December 2011, <http://www.computing.co.uk/ctg/news/2133724/jp-morgan-goes-live-maxeler-supercomputer>. News article reporting JP Morgan’s deployment of Maxeler’s FPGA-based high performance supercomputer for risk evaluation. Archived here on 22/10/2013 <https://www.imperial.ac.uk/ref/webarchive/hyf>

[D] O. Mencer, E. Vynckier, J. Spooner, S. Girdlestone and O. Charlesworth, “Finding the right level of abstraction for minimizing operational expenditure”, Workshop on High Performance Computational finance, pp.13-18, 2011. <http://dx.doi.org/10.1145/2088256.2088262>. Publication by Scottish Widows and Maxeler on the effectiveness of their HPC on financial computing by Scottish Widows.

[E] S. Weston, J. Spooner, S. Racaniere and O. Mencer, “Rapid Computation of Value and Risk for Derivatives Portfolios”, Concurrency and Computation: Practice and Experience, 24(8): 880-894, 2012. <http://dx.doi.org/10.1002/cpe.1778>. Publication by JP Morgan and Maxeler on using FPGA-based HPC in 2009 and 2010 on risk computation with 30x performance acceleration.

[F] “Maxeler Technologies Sells 20 Percent Stake in Company to J.P. Morgan”, March 31, 2011. <http://www.hpcwire.com/hpcwire/2011-03-31/maxeler-technologies-sells-20-percent-stake-in-company-to-j-p-morgan.html>. News article reporting JP Morgan taking a 20% stake in Maxeler. Archived on 22/10/2013 <https://www.imperial.ac.uk/ref/webarchive/jyf>

[G] Managing Director and Global Head of Applied Analytics Group at J. P. Morgan, letter stating the impact of our reconfigurable research on J. P. Morgan’s business with quantitative estimates.

[H] T. Nemeth, J. Stefani, W. Liu, R. Dimond, O. Pell and R. Ergas, “An Implementation of the Acoustic Wave Equation on FPGAs”, 78th Society of Exploration Geophysicists (SEG) Annual Meeting, November 2008. <http://www.maxeler.com/media/documents/MaxelerSummaryAcousticWaveEquation.pdf>. Archived [here](#) on 22/10/2013

[I] CEO of BlueBee Technologies, stating the impact of our research on their products.

[J] News and views from the HL Steam team, “HL Steam teams up with BlueBee for hardware acceleration”, 16 November 2012.

<http://hlsteam.com/hl-steam-teams-up-with-bluebee-for-hardware-acceleration> Archived on 22/10/2013 at <https://www.imperial.ac.uk/ref/webarchive/kyf>