Institution: University of Bristol		
Unit of Assessment: 11) Computer Science and Informatics		
Title of case study: Enabling Voice Recognition Everywhere		
Period when the underpinning research was undertaken: 2000-2007		
Details of staff conducting the underpinning research from the submitting unit:		
Name(s): David May	**Role(s) (e.g. job title):** Professor of Computer Science	**Period(s) employed by submitting HEI:** Sept 1995 – Sept 2020
Period when the claimed impact occurred: 2014 – 31 July 2020		
Is this case study continued from a case study submitted in 2014? Y		

1. Summary of the impact

The original impact of Professor May's research, and subject of the REF2014 return, was the applicability of a novel computer architecture in audio markets, such as audio over Ethernet. This work led to the formation of XMOS Ltd, a University of Bristol spin out, attracting over USD30 million in investment since 2014. During the REF2021 period, this same architectural approach of low-cost, predictable performance and real-time processing has proven to be ideal for enabling the rapidly expanding market in voice processing in edge devices, such as smart speakers in the home. XMOS' processors now appear in products supporting ubiquitous voice-enabled technologies such as Alexa from Amazon. XMOS' revenue has subsequently grown substantially, from the GBP4 million p.a. reported in REF2014 to over GBP10M p.a. in 2018, with the company growing in size by 40% to support these new markets.

2. Underpinning research

The research leading to the creation of XMOS was undertaken at Bristol originally during the REF2014 return period. It involved innovations in the architecture of processor cores for use in processor arrays on chips, in modules and in systems (along with the associated programming concepts). In particular, it developed novel processor architectures combining a unique combination of deterministic performance, multi-threaded processors, interprocessor interconnect and energy-efficient, event-driven processing.

Between 2000 and 2001, May's research focused on multi-threaded processors in which each thread could deliver deterministic performance, a unique ability at that time [3.1, 3.2]. This research included instruction set design and implementation techniques for thread initiation, synchronisation and termination, along with thread register addressing and pipeline organisation, with a goal of delivering predictable, deterministic performance, a significant advantage over the current state of the art in embedded processors. In 2000, May also demonstrated compiler technology for optimising concurrent programs and for parallelising sequential programs by distributing execution of a sequential program among several threads [3.2], a key supporting technology to deliver the benefits of the processor he'd been designing.

Further work by May in 2000 and 2001 investigated how the communicating process architecture – which was pioneered in the Inmos transputer, also developed by May and his University of Bristol team, including Dr Muller (Reader in Computer Science at Bristol from 1993-2008, now Principal Technologist at XMOS), could enable new classes of products [3.3, 3.5]. This work developed important concepts for multi-threading and caching techniques, along with mobile channels and mobile processes, which underpin the resulting XMOS impact.

Also in 2001, May explored the design of instruction sets to support highly parallel processors [3.4]. The XMOS architecture is based on multithreaded processors, drawing directly on [3.1] and [3.2] to provide instruction set and hardware support for deterministic execution, multi-threading, communication and synchronisation. Many aspects of the instruction set design, such as the support for position-independent object code to support process mobility and the use of system-wide channel addresses, were based on work described in [3.3]. Important techniques

for reducing the instruction set to the minimum needed to support high-level programming languages while also aiding the drive towards predictable, deterministic performance, were developed [3.4].

Prof. May has filed fifteen patents on the core technology which uniquely combines multi-threading for energy efficiency and performance while also delivering predictability and deterministic execution [3.6]. This novel combination made the XMOS architecture ideal for real-time applications such as audio over ethernet and more recently voice recognition, all at hitherto unreachable levels of energy efficiency and cost competitiveness.

3. References to the research

3.1 Towner D and **May D** (2001). The 'Uniform Heterogeneous Multi-threaded' Processor Architecture, in *Communicating Process Architectures*. Chalmers A, Mirmehdi M, and Muller H, (eds.). pp.103-116. ISBN 1 58603 202 [Available on request]

3.2 Sen S, Muller H, and **May D** (2000). Synchronisation in a Multithreaded Processor, in *Communicating Process Architectures*. Welch P & Bakkers A, (eds.). pp.137-144. ISBN 1 58603 077 9 [Available on request]

3.3 **May D** (2000). The Transputer Revisited, in *Millennial Perspectives in Computer Science*, Palgrave, pp. 215-228. ISBN 10: 0333922301 / ISBN 13: 9780333922309 [Available on request]

3.4 **May D** (2001), *Picochip Instruction Set Proposal*. CSTR-06-012, [Available on request]

3.5 Watt D and **May D** (2001). A Language and Processor for Unifying System-on-Chip Design, in *Communicating Process Architectures*. Chalmers A, Mirmehdi M, & Muller H (eds), IOS Press, pp.167-178. ISBN 1 58603 202 X [Available on request]

3.6 15 patents granted to Prof. May on the XMOS core technology, including threading, scheduling, instruction set encodings, event control and message handling. Patent numbers: US 7613909 (filed 17 April 2007), US7617386, US 7676653, US 7948060, US 7958333, US 7962717, US 8139601, US 8185719, US 8219789, US 8224884, US 8347312, US 8898438, US EP 2499568, US 8966488, US 9367321

4. Details of the impact

The University of Bristol spin-out XMOS Ltd. was created by Professor David May in 2005. Its headquarters are in Bristol, with offices in the US (New Hampshire) and Hong Kong.

The subject of the previous REF was the first generation of XMOS' flagship product, the xCORE, the architecture for which is based on the University of Bristol research published in [3.1-3.5]. The IP that was spun out was a unique architecture designed to be intrinsically suited for real-time systems. The processors make it easier to design complex embedded systems for a range of consumer, audio, industrial and automotive applications. The same architecture from the previous REF has turned out to be ideal for new markets in voice processing, artificial intelligence (AI) and machine learning (ML), and so, in February 2020, XMOS announced a new product, the world's lowest-cost AI microcontroller [5.5]. It offers electronics manufacturers uniquely short development cycles, an agile build process, energy efficient processors, and the ability to correct faults quickly.

Economic growth of the XMOS spin-out / launch of xCORE-200

XMOS has seen significant growth during this REF period. In 2015, the company launched the second generation of xCORE, the xCORE-200. This was followed by the launch of xCORE-AUDIO and xCORE-VOICE in April 2016, and xCORE VocalFusion in June 2017.

xCORE-200 retained the first generation's core architecture, and thus applicability for voice processing in edge devices (e.g. smart speakers, voice-enabled Internet of Things, etc.), while also adding new features to improve performance, such as the addition of a dual issue mode, which doubled the peak performance, quadrupling the memory on the chip, and doubling the number of cores, all of which enabled superior Digital Signal Processing (DSP) capabilities and provided the computational power for designing far-field microphones. Enabling the design of far-field microphones is a key advance, as the modality of interfacing with computers and appliances is changing from being button-driven (e.g. TV remote control, buttons on a cooker-hood) to voice-driven. Capturing voice in a cooker-hood requires picking out the voice from the noise of the fan. Capturing voice in a TV requires picking out the voice over the sound of the programme that is playing. The DSP associated with these two computations is significant and is something that the second generation XMOS device is capable of performing [5.1] by exploiting the innovative multi-threaded processors, instruction sets and deterministic performance developed in [3.1- 3.5].

xCORE-200 has enabled XMOS to grow significantly from the last REF period, from 50 to 70 staff, and in annual revenues from GBP4 million in 2014 to over GBP10M in 2020 [5.1]. XMOS has sold more than 10 million of these new devices to date, mostly in the prosumer audio market, with audio brands such as FocusRite, Sony, Meridian and Native Instruments forming a significant part of the number of devices sold [5.1].

Enabling voice interaction in new products across the global Voice and Edge-AI market

The high performance, low-cost and deterministic performance of the original XMOS architecture has proven ideal to tackle the emerging "voice recognition at the edge" market. Chips and software developed by XMOS are enabling voice interaction across a wide range of products. AudioExpress (2017) stated that: "*The new XMOS VocalFusion 4-Mic Dev Kit features a compact four-microphone linear 100mm array. The captured voice signals are crystal clear even in noisy environments, enabling commands to be accurately captured from across the room for processing by the Alexa cloud-based speech recognition system. XMOS is the only solution to date that packs audio digital signal processing (DSP) and programmable I/O processing and control software into a single chip, enabling customers to significantly reduce their bill of materials, flexibly tailor their designs, and reduce costs*." [5.2]

The technology allows XMOS to offer their customers an integrated solution at a competitive market point. Meridian, a UK manufacturer of high-end audio products, has stated that: "*Using XMOS we were able to replace three separate development systems, saving us over 6 months in development time. All our new products now use xCORE.*" [5.1]. Meridian's products which have adopted XMOS include Meridian Explorer, Director and the Control 15 Media Server.

In 2017, XMOS' VocalFusion technology was selected by Pegatron Corporation for its new smart voice assistant, Martina. The technology delivers clear voice capture, the ability to connect to Google Cloud services and language recognition in four languages. [5.3]

In 2018, Tymphany, a global premier audio ODM, announced the latest project in building a new Alexa Built-In soundbar incorporating XMOS' VocalFusion far-field voice processor. The soundbar connects with the Amazon Alexa Voice Service (AVS) via WiFi and has an excellent far-field voice capture and barge-in performance from across the room – even when the soundbar is playing content at high volume and the commands are spoken softly. [5.4] Tymphany's Senior Director of Technology Management states: "*We are proud of our team's efforts to realise this complex, groundbreaking and great performing soundbar reference design. We would not have been able to achieve this milestone without the awesome support of the Amazon Alexa Voice Service and XMOS development teams*". [5.4]

In 2019, Amazon certified the use of the XMOS VocalFusion kits for Alexa-enabled devices, one of very few products who have this qualification; see Figure 1. [5.6] The Director of the Alexa Voice Service at Amazon has said: *"XMOS provides the thin-form factor and high performance needed […] to add Alexa in new ways to stereo home entertainment and wall-mounted AV*

equipment." [5.7]. The Chief Evangelist for Alexa at Amazon, went on to say about XMOS "*Having partners like you build hardware means we have a way for everybody to build a consumer device that you can talk to.*"

Figure 1 - XMOS Architecture qualified for use in Alexa [5.5]

Other XMOS customers include:

- Freebox, a large French TV operator who have the XVF3500 stereo-AEC voice processor built into their Delta product [5.5]
- Pillo Health, where XMOS voice technology underpins the 24/7 healthcare robot Pillo. The device delivers therapy reminders, health information, stores pills and reminds users of their scheduled medication doses. Emanuele Baglini, CTO of Pillo Health said: *"While designing our device, we were looking for the best solution to improve the voice capture accuracy to give our users a great experience. We tested the XMOS voice processor against other solutions on the market, and it outperformed them all for clarity, accuracy, distance, and direction.*" [5.5]
- Murata, where XMOS XVF3000 voice processor is built-in to their NAONA Array Mic unit, delivering the capability to capture speech and the direction of speech from a distance [5.5]
- Sharp, the Japan-based multinational who are using the XMOS VocalFusion XVF3500 voice-capture technology as part of their new Windows-based video conferencing and remote collaboration solution; see Figure 2 [5.5]
- Orange, where the XMOS VocalFusion XVF3000 enables their Djingo smart speaker [5.5]
- FiiO, the Chinese company producing high-quality music devices, includes XMOS in a range of their products [see, for example, 5.8]

Figure 2 - XMOS far-field voice-capture technology used in the Windows collaboration display, by Japan-based Sharp [5.7]

The third generation of XMOS' products, still based on the same underpinning architecture, went into volume production in the second half of 2020. XMOS' CTO has stated that "...*the technology that was spun out of the University of Bristol 15 years ago has turned out to be so generic that it has propelled XMOS into Voice and Edge-AI markets that nobody had foreseen in 2005. We are looking forward to what the next decade will bring us*." [5.1]

5. Sources to corroborate the impact

5.1 XMOS – Corroborating statement, CTO (June 2020)

5.2 AudioXpress – Article (October 2017), *XMOS Delivers Amazon Alexa Voice Service Development Kit with Linear Mic Array for Far-Field Voice Capture*, [Accessed 5 February 2021]

5.3 AudioXpress – Article (December 2017), *XMOS Enables Far-Field Voice Technology for Pegatron Martina Smart Speaker Interface to Google Cloud Services*, [Accessed 7 May 2020]

5.4 Tymphany and XMOS press release (August 2018), *Tymphany and XMOS to showcase new soundbar at IFA with Alexa built-in* [Accessed 19 October 2020]

5.5 XMOS website and press releases, xmos.ai, Nov 2018 – Feb 2020 [Accessed 7 May 2020]:
[i] *XMOS 2-mic dev kit for Alexa Voice Service qualified by Amazon*;
[ii] *XMOS celebrates the launch of Freebox Delta*;
[iii] *XMOS wins at Made In The South West Awards 2018*;
[iv] *XMOS helps create a sensor for Murata's sensor data platform*;
[v] *Introducing Windows collaboration display by Sharp*;
[vi] *"Ok Orange!" XMOS helps bring the Orange Djingo smart speaker to life*;
[vii] *XMOS announces world's lowest cost, most flexible AI processor*

5.6 Development kits for Alexa Voice Service (2019), https://developer.amazon.com/en-US/alexa/alexa-voice-service/dev-kits [Accessed 7 May 2020]

5.7 Amazon Alexa (February 2018) – Blog, *New XMOS Development Kit for Integrating Alexa into Slim Profile and Wall-Mounted Devices* [Accessed 7 May 2020]

5.8 FiiO website, https://www.fiio.com/q5s [Accessed 18 August 2020]