# UDRC Themed Meeting on Algorithm Implementation and Low SWAP Challenges

As part of the UDRC phase III, a themed meeting Algorithm Implementation and Low SWAP Challenges will be held on Wednesday 30th November. This will be an in-person event. This event is intended for academic researchers, industrial partners, and Dstl staff to learn about and discuss current trends in this area. The program will consist of a series of talks from academia and defence industry along with panel discussions.

**Timings:** Wednesday 30th November 2022 (coffee at 9am and start at 9:20am)
**Location:** James Watt Centre, Craig Room, Heriot-Watt University
**Registration:** https://edinburgh.onlinesurveys.ac.uk/attendance-udrc-themed-meeting-algorithm-implementation-an

**Organisers**: Dr Joao Mota (j.mota@hw.ac.uk); Prof Andy Wallace (a.m.wallace@hw.ac.uk)

| Time | Topic | Name |
|---|---|---|
| **09:20 – 11:00** | **Session 1 – Approximation Guarantees, Chair & Introduction: Andy Wallace** | |
| 09:30-10:00 | A compiler for sound floating-point computations | Joao Rivera, ETH Zurich |
| 10:00-10:30 | Why does scientific high performance computing need FPGAs? | Nick Brown, EPCC |
| 10:30-11:00 | Reconfigurable approximate accelerators for signal processing on resource constrained systems: from algorithms to real-time implementation | Yun Wu, Heriot-Watt University |
| **11:00 – 11:30** | **Break** | |
| **11:30 – 12:00** | **Session 2 – Efficient AI, Chair: Joao Mota** | |
| 11:30-12:00 | The challenges of AI and multi-camera imaging - an MBDA perspective | Nikki Easton/Dean Goff, MBDA Systems |
| **12:00 – 13:30** | **Lunch Break** | |
| **13:30 – 14:30** | **Session 3 – Design and Algorithm Implementation, Chair: Mathini Sellathurai** | |
| 13:30 – 14:00 | Low-SWAP real-time EMG signal processing for human-computer interaction | John McAllister, Queen's University Belfast |
| 14:00 – 14:30 | Better performance through resource-elastic dynamic stream processing on FPGAs | Dirk Koch, Heidelberg University Germany / University of Manchester |
| **14:30 – 15:00** | **Break** | |
| **15:00 – 16:00** | **Session 4 – SWAP Communication and Sensing, Chair: Bernie Mulgrew** | |
| 15:00 – 15:30 | Decoupling and parallelisation for broadband multichannel problems | Stephan Weiss, University of Strathclyde |
| 15:30 – 16:00 | Processing and SWAP challenges for digital radar | David Greig/Hannah Durnall, Leonardo |
| **16:00 – 16:30** | **Closing remarks / discussion** | |

*The University Defence Research Collaboration in Signal Processing in the information Age is funded by EPSRC and Dstl.*

Engineering and Physical Sciences Research Council

# Abstracts

## A Compiler for sound floating-point computations, Joao Rivera, ETH Zurich

Floating-point arithmetic is widely used in scientific and engineering applications but is unsound, i.e., there is no guarantee on the accuracy obtained. When working with floating-point, developers often use it as a proxy for real arithmetic and expect close to exact results. Unfortunately, while often true, the result may also deviate due to round-off errors. Their non-intuitive nature makes it hard to predict when and how these errors accumulate to the extent that they invalidate the final result. In this talk, we discuss an automatic approach to soundly bound such uncertainties. We present a source-to-source compiler that translates a C program performing floating-point computations to an equivalent C program with soundness guarantees, i.e., the generated program yields a precision certificate on the number of correct bits in the result. We will briefly discuss the design and implementation of our compiler, followed by techniques used to improve the performance and accuracy of the result.

## Why does scientific high performance computing need FPGAs?, Nick Brown, EPCC

Scientists and engineers are continually demanding the ability to generate more detailed results at reduced time to solution. Little wonder then that, in an attempt to meet such ambition in an energy efficient manner, the future role that novel hardware architectures, such as Field Programmable Gate Arrays (FPGAs) can play in High Performance Computing (HPC), is of great interest. However, to date, FPGAs have struggled to gain traction in HPC and whilst there were some historical efforts, at the time the technology was not really mature enough. However, times change and the past few years have seen very significant advances in both the hardware and software ecosystems, potentially making FPGAs a far more competitive and accessible acceleration target. In this talk I will describe the major advantages that we consider in leveraging FPGAs for scientific workloads and, using a number of real-world case studies as examples, will explore the benefits that the technology has delivered.

## Reconfigurable approximate accelerators for signal processing on resource constrained systems: from algorithms to real-time implementation, Yun Wu, Heriot-Watt University

Modern signal processing systems adopt increasingly complex algorithms with sophisticated functionalities to meet remarkable performance requirements. Due to the volume of algorithmic arithmetic operations, it is challenge for digital signal processors to attain the computational capacity with real-time processing, while maintaining the system resources constraining threshold, such as space, power, and cost. It is increasingly important to employ the emerging contexts of approximate computing in computer science research to structure signal processing algorithms and systems, by trading off the accuracy of results with the utilization of resources for real-time implementations. In this talk, we are going to introduce our latest research targeting such challenge in signal processing domain, especially those applications of solving convex optimization problems. By introducing our approximate computing approach, we are going to present our top-down design flow for reconfigurable accelerator generation, from algorithm to real-time implementation, with insight of algorithmic and design efficiencies. Furthermore, we are going to discuss some future research directions from our on-going works.

*The University Defence Research Collaboration in Signal Processing in the information Age is funded by EPSRC and Dstl.*

## The challenges of AI and multi-camera imaging - an MBDA perspective, Nikki Easton/Gary Matson/Dean Goff, MBDA Systems

The relatively new fields of AI and multi-camera imaging have been demonstrated to show enormous promise; AI in the huge leap forward in computer vision algorithmic performance, and multi-camera imaging with the ability to locate cameras off boresight, remove gimbals, image in multiple wavebands and improve both signal-to-noise ratio and angular resolution by fusing the images from multiple sources. Both fields however put significant pressure on computational hardware to ensure they can operate at high frame rates. This presentation provides an insight into the early development work being undertaken at MBDA UK to enable these technologies to be exploited in future products, including investigations and challenges encountered when working with the Texas Instruments Jacinto to accelerate custom AI architectures.

## Low-SWAP real-time EMG signal processing for human-computer interaction, John McAllister, Queen's University Belfast

Electromyograms (EMGs) are electrical signals emitted by muscles during movement and have long been targeted as a medium by which human movement can be inferred and used to control computers powering bionic limbs or interact with virtual or augmented environments. In general, these initiatives have resulted in algorithms, which are extremely computationally complex, require sensitive installation or do not operate in real-time using battery-powered devices. However, a new generation of consumer devices exploiting multi-channel randomised EMG sensing proposes a very different sensing and processing modality and a very different set of processing challenges and opportunities. This talk outlines how these approaches lend themselves not only to real-time processing on low-cost processors, but also open up very many degrees of freedom via which to trade SWAP with accuracy.

## Better performance through resource-elastic dynamic stream processing on FPGAs, Dirk Koch, Heidelberg University Germany / University of Manchester

This talk will show how pipelines built of composable operator modules can be stitched together using partial reconfiguration on FPGAs. This allows it to build optimized accelerator pipelines for problems that are only known at runtime. Moreover, the approach allows us to maximize the amount of work performed per unit I/O for improved performance and energy efficiency. The proposed system provides a full ecosystem comprising supporting IP, compile scripts for implementing stitchable accelerator modules, a bitstream manipulation tool to relocate modules, and a runtime system. The latter can compile user requests into operator modules, orchestrate the configuration and initialization of modules, and drives the execution of the hardware. The approach will be demonstrated on SQL database acceleration where an FPGA is composing optimized query accelerators on-the-fly. In this system, partial FPGA configuration is inferred transparently for the user just by requesting some SQL queries.

## Decoupling and parallelisation for broadband multichannel problems, Stephan Weiss, University of Strathclyde

High-dimensional data is often preprocessed such that the processing task can be split into a number of separate, independent sub-tasks. Typical examples include the discrete Fourier transform (DFT) or filter banks, which aim to address a processing challenge in separate frequency bins or subbands. Subsequently, data across frequency bins is often assumed to be mutually independent, admitting

*The University Defence Research Collaboration in Signal Processing in the information Age is funded by EPSRC and Dstl.*

parallel implementations. For broadband signals, spectral coherence implies that the independence assumption of frequency bins is generally unjustified. In this presentation, we will focus on the analytic eigenvalue decomposition (EVD) as a broadband equivalent to the standard EVD or Karhunen-Loeve transform. The analytic EVD is capable of strongly decorrelating signals via a lossless filter bank operation, such that the independence assumption of the emerging subbands is indeed satisfied. We will justify this via the coding gain of the created subband signals, and comment on the computational complexity of the transform, and the complexity of the created subtasks for some example applications in the acoustic/RF domains.

## Processing and SWAP challenges for digital radar, David Greig/Hannah Durnall, Leonardo

Digital radars have many advantages in terms of multi-functionality and multiple beam-forming. This is also of intrinsic value both in terms of cognitive processing during a mission and analysis post mission. Increased capability always comes with a cost, in this case processing power, data storage and data retrieval within the constraints on size weight and power (SWAP) for an airborne sensor. In this talk, we shall consider how the data and processing challenges change as radars become increasingly digital and the challenges on processing requirements this will bring including low latency as well as limitations on SWAP. Consideration will be given to possible processor architectures, novel algorithmic methods and the application of machine learning.

*The University Defence Research Collaboration in Signal Processing in the information Age is funded by EPSRC and Dstl.*