



University Defence Research Collaboration (UDRC) Signal Processing in a Networked Battlespace

E_WP6: Efficient Computation of Complex Signal Processing Algorithms
WP Leader: Prof. John Thompson, University of Edinburgh
Researcher: Calum Blair, University of Edinburgh

Introduction: Signal processing algorithms are becoming more sophisticated. Sensors are becoming more ubiquitous, and processors more powerful. Massively parallel processors such as General-Purpose Graphics Processing Units (GPUs) and Field-Programmable Gate Arrays (FPGAs) allow fast, often real-time processing of certain algorithms. Any deployment on a device will have a fixed **Size, Weight and Power (SWaP)** which affects algorithm runtime and, indirectly, complexity. **Design Space Exploration** involves choosing a device and investigating algorithm modifications to fit the platforms they run on.



Algorithms must be modified to suit the device they are deployed onto. However, knowledge of device behaviour can be used to change algorithms at the design stage.



Any implementation on a device is a *tradeoff* between **algorithm effectiveness or accuracy**, and **number of computations required**, which directly impacts runtime and **SWaP**. Separately, **Complexity and Ease of Parallelization** must be considered (i.e. how *embarrassingly parallelizable* is the algorithm?)


Challenges: This Work Package considers algorithms and algorithmic improvements which apply to areas throughout signal processing. We seek to identify algorithms which deploy onto hardware easily or with limited modification to produce **improvements in SWaP requirements** over the state of the art. This can include reductions in the number of computations required, with **minimal performance loss**. We are addressing the following challenges from **dstl**:

27: Accreditable Machine Learning or Data-Driven Techniques


Reliable testing and accreditation of **machine learning algorithms** is problematic, particularly when the test environment for an algorithm differs significantly from the training dataset – known as the *covariate shift* problem. A method for calculating the reliability of a machine learning algorithm in a wide variety of scenarios is required.

29: Reducing Size, Weight and Power Requirements through Efficient Processing


FPGAs and GPUs on embedded platforms allow greatly increased computations while reducing power consumption. Parallelsed processors can be combined with techniques which **reduce computational complexity of an algorithm** and used to accelerate algorithms such as **Support Vector Machines** or **Gaussian Mixture Models**. The end goal here is *accurate algorithms running on an embedded device in real-time*. As we work on techniques with multiple applications across the signal processing domain, we will interact with and support the other work packages in the Edinburgh Consortium as shown below:



Approaches: Our current approaches to **C27** and **C29** are below, and are investigated as part of *E_WP6.1: Efficient Parallelisation of Sensing Processing*. Previous work¹ involved parallelisation of multiple algorithms in a heterogeneous system, looking at tradeoffs between power and time on three platforms (FPGA, GPU, CPU):



Introspective Classification builds on this by evaluating the *uncertainty or entropy* of any classification or detection, as well as returning a confidence measure. We focus on **pedestrian detection**, as techniques here extend well to **other object detection tasks**.

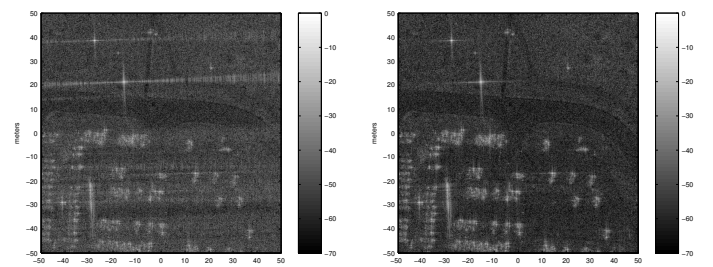


Multi-object classification with cars (yellow) and pedestrians (blue).

False detections of pedestrians (blue) interferes with higher-level tasks such as anomaly detection.

By reducing false alarms at the detection stage and by providing downstream trackers with *improved probabilistic detections*, overall performance can be improved. Uncertainty information also improves the ability to flag *unseen object classes not present in the training set*².

Sparse Imaging for Synthetic Aperture Radar (SAR): RA Shaun Kelly developed an iterative sparse algorithm for reconstructing images from undersampled SAR data during UDRC Phase 1 work³. Current runtimes are prohibitive on both desktop and mobile platforms so Shaun is parallelising and porting this to **GPU**.



This is closely related to GPU-accelerated hybrid simulation for **Synthetic Aperture Sonar (SAS)**, which we will investigate in conjunction with *E_WP3*.

Future Work: As part of *E_WP6.2: Implementation of Distributed Signal Processing Algorithms* we will use a **mobile sensing platform** in combination with other static sensors to investigate **computation/communication tradeoffs** in a distributed network. I.e. in such a network with heterogeneous nodes, what is the optimal distribution between per-node computation and communication and how does this change with time, processing required and number of targets? Following this we will investigate *E_WP6.3: Algorithm & Compute Resource Management*; given finite computational resources at a node with one or more sensors, which sensor or region should be prioritised most, and by how much? We will again consider how this changes with varying target numbers. Collaboration with *E_WP5* is expected here.

Conclusion: The broad theme of this work involves reductions in algorithm complexity, and reductions in detection false alarm rate. Both these approaches allow improvements in various areas throughout the domain of signal processing.

¹ Blair, C.G. & Robertson, N.M., "Event-Driven Dynamic Platform Selection for Power-Aware Real-Time Anomaly Detection in Video". In Proc. VISAPP 2014 (To appear).
² Grimmel, H., Paul, R., Triebel, R., Posner, I., "Knowing When We Don't Know: Introspective Classification for Mission-Critical Decision Making". In Proc. ICRA 2013.
³ Kelly, S.I., Du, C., Rilling, G., Davies, M.E., "Advanced image formation and processing of partial synthetic aperture radar data". IET Signal Processing, 6(5), 2012.

