EPSRC Centre for Doctoral Training in High Performance Embedded and Distributed Systems (HiPEDS)

Imperial College London
Department of Computing
Department of Electrical and Electronic Engineering

Advisory Board Meeting 5th of December 2014

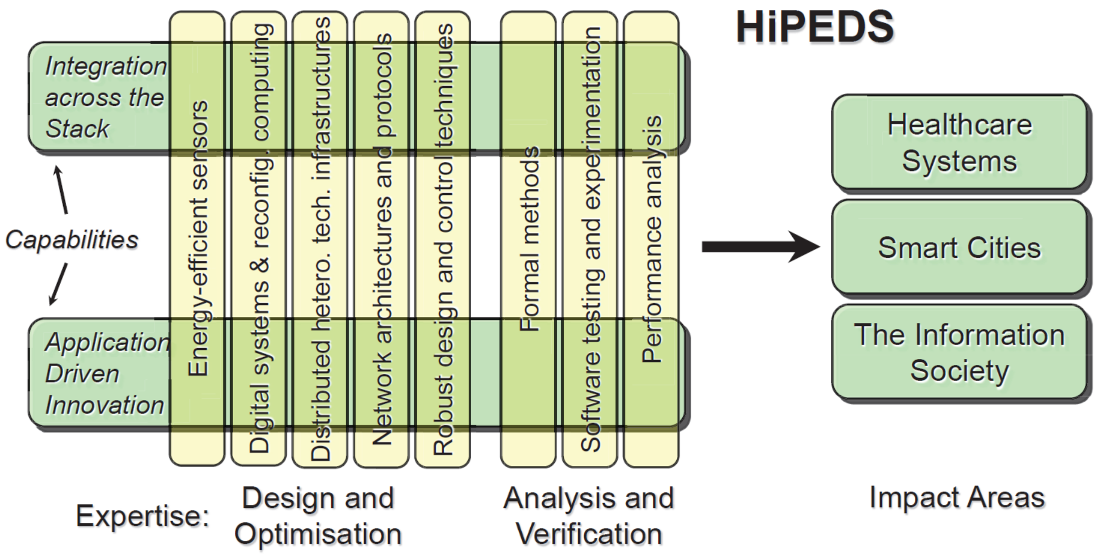

Student Project Narratives 2014

Contents

1. **Ira Ktena (Daniel Rueckert, DoC)**

Computational Analysis of Brain Connectivity during Early Development

Our brain composes a high performance distributed system in itself. The average human brain has about 100 billion neurons that constitute the central nervous system, and each of them may be connected to up to 1,000 other neurons, yielding almost 100 trillion synaptic connections. Nowadays, modern magnetic resonance imaging (MRI), and neurophysiological techniques have provided efficient and non-invasive ways to map the patterns of structural and functional connectivity of the human brain. The advancement of acquisition techniques at an unprecedented resolution contributes significantly to this endeavor and opens up new horizons for the investigation of the human connectome, which is one of the greatest scientific challenges of the 21st century. The main objective of this research effort is to elucidate the neural pathways that underlie brain function and behavior, and decipher this amazingly complex wiring diagram. Among other things, this procedure will potentially reveal much about what makes every person different from all others and to which extent this network structure is attributable to genetic factors. Network analysis provides an ample set of tools for exploring the association between brain functional deficits and the underlying structural disruption which is related to brain disorders and the dynamics of these networks can potentially shed light on disorders like autism, whose mechanisms are yet to be explored.

2. **Antoine Toisoul (Abhijeet Ghosh, DoC)**

Acquisition and Simulation of Wave Effects in Material Appearance

The project will involve developing acquisition systems for laboratory and on-site measurements of wave optics effects in appearance of real world materials as well as GPGPU based computations for realistic rendering and inverse rendering simulations for realistic computer graphics applications. For the acquisition component, the aim will be to develop controlled programmable illumination systems for laboratory measurements including networked LED lights, wireless or LED flashes, and controlled measurement with networked machine vision or DSLR camera systems. For on-site acquisition, the focus will be on developing acquisition techniques using mobile devices and other portable sensors. The project will also investigate acceleration of simulations of wave effects in material reflectance and scattering using GPU and GPGPU computations for rendering as well as inverse rendering. Thus, the project fits into the two cross-layer research themes of design and optimization for the acquisition and rendering components and analysis and verification for the inverse rendering and appearance estimation components. The primary application area of the project is expected to be the Information Society, with possible applications also in Healthcare Systems.

3. Adrian Rapeaux (Tim Constandinou, EEE)

Embedded Neural Microsystems

This project is located at the intersection between embedded microsystems and resource-efficient implants and sensors. Within the HiPEDS CDT, it is a Design and Optimisation project with focus on physical volume and energy efficiency. The aim is to create an implantable and compact neural interface as an autonomous embedded system-on-chip (SOC). Using novel stimulation techniques together with FDA-approved electrode designs, major drawbacks to existing nerve implant technologies can be addressed and provide the next generation platform for neural implants. There is currently a strong commercial drive for targeted nerve stimulation technologies (e.g. GSK electroceuticals initiative). Amongst existing products, several targets for improvement have been identified, such as fragile power leads to external battery cases, low stimulation selectivity, and insufficient system communication capabilities. Implementing an embedded, autonomous neural interface as a SOC would address these issues and pave the way for better implant technologies. The results that can be achieved with this approach are particularly interesting with HiPEDS. The design specifications of an implantable, compact, autonomous system fall squarely into the field's area of expertise. The tools needed to create a full-custom SoC are available, and the applications of such a system within healthcare are numerous and of interest to both industry and research.

4. Salvatore Dipeitro (Giuliano Casale, DoC)

Dependability of distributed cloud systems

Salvatore Dipietro's research will focus on dependability of distributed cloud systems. Cloud applications are more scalable and cost-effective that traditional data centre applications thanks to the resource elasticity offered by cloud infrastructures. However, since these applications run in a shared execution environment, they are exposed to an increased number of threats to their availability, security, and reliability compared to traditional data centre applications. For example, resource contention by other tenants on a cloud platform can affect the amount of compute, memory and I/O resources that are allocated to the application, as a result or malicious or unintentional behaviours of other tenants. The goal of this research will be to define techniques to expose the presence of such threats through the analysis of operational data of the application coming from log-files and advanced monitoring. Workload analysis techniques will be used to detect outliers in the application behaviour both in centralised and distributed scenarios and to automatically build workload profiles that can support inference of dependability threats. This research will be at the interface between systems research, data mining and performance evaluation. The results will foster the definition of novel techniques to enhance the dependability of distributed cloud applications.

5. Christian Priebe (Peter Pietzuch, DoC)

EP-SRC CloudSafetyNet

The project looks at the timely topic of security mechanisms for cloud environments, it aligns naturally with the goals of the HiPEDS CDT. Christian's PhD would add a security angle to the topics covered by HiPEDS. The ultimate goal of his PhD work, according to the CloudSafetyNet project, would be to design and evaluate a secure cloud computing stack, which includes hypervisor-based mechanisms for the tracking of sensitive data through the cloud infrastructure.

6. Konstantinos Boikos (Christos Bouganis, EEE)

Design of embedded computer vision systems

The student's research focuses on the design of embedded computer vision systems i.e. vision processors. Current state-of-the-art vision processors have been designed for the acceleration of specific tasks such as navigation, surveillance or augmented reality support. As such, they usually have a fixed architecture, which does not perform well when it is applied to a task that has not been tuned for, or when some of the environmental assumptions are not met. The research will consider the application of reconfigurable computing to embedded computer vision systems in order to provide flexibility when the vision processor is applied to various tasks. The envisioned flexibility has the potential to lead to a vision processor that can adapt its architecture to the needs of the application and the environment, resulting in better performance and power characteristics. In addition, the intention is to research machine learning techniques that would enable such on-line adaptation, but under strict performance and power budgets. The project is well suited within HiPEDS as the envisioned processor needs to be of high-performance and low-power consumption, taking advantage of the inherent characteristics of vision applications in terms of approximations in computations, and being able to adapt to its surroundings.

9. Eric Seckler (Alex Wolf, DoC)

Architecting distributed software

Eric is interested in distributed systems, generally, and in the specific problem of architecting distributed software according to novel design criteria. My group is facing several such problems right now, mainly in the context of cloud data centres. We are exploring the use of heterogeneous hardware and networking technologies (GPUs, FPGAs, ASICs, network middle-boxes, SSDs, and the like) to serve the disruptive needs of non-traditional cloud applications. In that context, we are finding that applications require new and innovative structuring principles to allow them to be flexibly and optimally managed across heterogeneous platforms. This is where Eric's interests can come into play: developing new architectural principles, and the programming tools to support them, driven by novel design criteria. This problem is well aligned with HiPEDS, since it requires knowledge and skills across an especially wide range of computing disciplines. Eric's PhD would need to touch upon nearly every one of the crosscutting areas of expertise.

8. Panagiotis Garefalakis (Peter Pietzuch, DoC)

Consistency models in scalable key/value stores

Panagiotis is interested in distributed computing, in particular how to design scalable, robust and secure distributed systems. An obvious direction for is PhD work would be to take his work on consistency models in scalable key/value stores and build further on top of this research. Such a direction would align well with the ICT-Labs-funded MC-Data project in my group, which investigates issues of data replication and consistency in cloud infrastructure that employ multiple data centres interconnected through wide-area network links. Such a topic would align with the focus of HiPEDS on distributed systems that can support real-world applications in a number of different domain areas.

9. Nicholaos Miscourides (Pantelis Georgiou, EEE)

Self-aware distributed sensing systems for robust and reliable chemical measurement.

This research will develop a novel methodology for CMOS based chemical sensing arrays that allows them to be fully self-aware, autonomous, accurate and always reliable when being used for continuous measurement. Chemical sensors when being used for continuous measurement of analysis of glucose or DNA are prone to degradation in performance, faults and also require human intervention for calibration. Moving forwards toward truly autonomous medical devices, we believe the sensing system should report an accurate value at all times, never fail in addition to avoiding any human intervention. A distributed architecture will be investigated as platform to achieve this consisting of multiple processing elements running in parallel and capable of measuring and making the system self aware by continually interrogating the sensors to assess their health and condition. Embedded algorithms will be developed to include fault tolerance always derive the most accurate sensing modality in addition to making the sensors fully autonomous through self-calibration. This will be a demonstration of the first truly self-aware and autonomous chemical sensing system.

10. Nicolas Moser (Pantelis Georgiou, EEE)

Distributed processing mechanisms for ISFET arrays to achieve in-pixel compensation, quantization and sensor enhancement.

This research aims to develop a new design methodology for CMOS based ISFET sensing arrays that will result in improved sensor performance for large-scale integration in applications such as DNA sequencing. ISFET sensors when used in large arrays are susceptible to trapped charge and temperature variation that can cause a significant variation in sensor output and limits their application. Additionally scaling to millions of sensors has limitations in data throughput. A new distributed architecture for ISFET sensing arrays with in-pixel processing and storing capability will be developed to compensate the non-idealities of the sensor and improve performance. The architecture will in a compact form implement an auto-calibration algorithm to reduce pixel variation and provide in-pixel quantization to significantly reduce the data bandwidth and be scalable across all CMOS processing nodes. A novel methodology for signal condition will also be applied to boost the sensors sensitivity and complete temperature independence will be achieved. It is anticipated that the application of this array will be for genotyping bacteria that is vital to combat the growing incidence of anti-microbial resistance.

11. Nadesh Ramanathan (George Constantinides, EEE)

Memory Optimization for Reconfigurable Computing

As the performance requirements of embedded systems increase, the fundamental bottleneck has become the memory subsystem, a familiar situation from general purpose and high-performance computing. However, embedded systems often run code that changes only irregularly, or sometimes not at all. This suggests that greater opportunities are available to optimise the memory subsystem of high performance embedded systems. For parallel memory accesses, a variety of programming models and memory guarantees are available, and these have differing energy consumption, cost, and performance implications. By taking a holistic approach to system design from source, we hope to be able to automatically derive highly optimised systems with radically improved energy / performance tradeoffs compared to traditional cache techniques.

12. Julian Sutherland (Philippa Gardner, DoC)

Reasoning about Distributed Algorithms

There has been a breakthrough in reasoning about concurrent programs. Previously, we had Owicki-Gries reasoning and rely-guarantee reasoning, which do not scale. Now we have reasoning about various forms of concurrent separation logics, which do scale: in particular, with TaDA, we are able to specify a concurrent set module and prove that an implementation using e.g. hash maps or concurrent B-trees is correct; and, with CoLoSL, we are able to reason about Dijkstra's famous token-ring mutual exclusion algorithm, a key algorithm in distributed systems. The aim of this project is to develop scalable reasoning about distributed systems, seeing what techniques apply from our concurrent reasoning. The first technical challenge will be to investigate how to reason about context-addressable networks, which can be viewed as distributed, Internet-scale hash maps. This project sits squarely within the HiPEDS cross-layer research theme of Analysis and Verification.

13. Emanuele Vespa (Paul Kelly, DoC)

Algorithms, architectures and domain-specific optimisations for dense SLAM

SLAM (simultaneous location and mapping) is the problem of building a 3D map of an environment, and tracking your path in that environment. SLAM is a key enabling technology for a wide range of applications in robotics, augmented reality, entertainment and telepresence. This project focuses on "dense" SLAM, where we reconstruct a model of the surfaces in the environment, as is required for motion planning, physics simulation and graphical rendering. The goal is to develop a "dense SLAM pipeline", analogous to the consensus that emerged in graphics rendering, which led to the growth of the GPU accelerator industry. This PhD project is focusing on designing fast algorithms for integrating camera data into the 3D map, and fast data structures for representing it in a way which allows sparsity to be exploited while aggregating information in a statistically efficient way. A key element of the research is to work in an autotuning framework that enables algorithmic variants to be evaluated for quality of result as well as energy and framerate.

14. James Targett (Wayne Luk, DoC)

Rapid simulation and evaluation of large dataflow designs

Current large reconfigurable designs often take hours or even days to compile, before they can be simulated. The proposed project aims to support rapid simulation and evaluation of large dataflow designs, based on a new generation of customisable instruction processors and the associated development tools. This approach will enable a shorter development cycle of novel designs targeting various applications including those within the HiPEDS impact areas such as healthcare, hence contributing to the two HiPEDS capabilities of (a) integration across the design and implementation stack, and (b) application-driven innovation.

15. Adrian Millea (Abbas Edalat, DoC)

Information Geometry and Deep Learning

Deep networks are parametric models, part of a rapidly advancing paradigm in the field of machine learning, called deep learning. Their success is due to the relative ease of deriving new models, the huge capacity of such deep models and also due to the increasing affordability of computing resources. Information geometry considers parametric models as points on high dimensional non-Euclidean manifolds where the metric on the manifold is given by the Fisher information, a way of measuring the amount of information that an observable random variable carries about an unknown parameter of a distribution that models the random varibale. Once we define such a manifold we are in the realm of differential geometry where we have a wide range of tools to work with. For example, the gradient descent iterations describe a parametric curve on this manifold, while keeping the relative entropy from one iteration to the next constant, gives rise to the natural gradient, which has been shown to perform better and converge faster than the usual gradient descent. We hope to gain insights into deep networks and their learning and then by making use of the set of tools available in information geometry to improve deep learning algorithms.