Imperial College London

Department of Computing
RA Symposium
Date: Thursday 11th June

Huxley Building Room 144

9:30 Welcome: Layal Hakim
9:35 - 9:50 Karen Hinxman *Postdoc Development Centre*

Session 1

9:50 - 10:10 Pavel Burovskiy
Custom hardware accelerators to the FEM: dealing with sparsity

10:10 - 10:30 Lawrence Mitchell
Symbolic numerical computing: solving equations by writing them down

10:30 - 10:50 Luigi Nardi and Zeeshan Zia
Multi-space multi-objective design-space exploration for 3D scene understanding using active learning

10:50 - 11:10 Nicholas Ng
Pabble, Session Types and dataflow programming

11:10 - 11:30 Jeremy Cohen
libhpc: A framework and tools for simplifying running of large-scale scientific HPC jobs for end-users

11:30 - 11:50 Poonam Yadav
Interested in converting a scientific problem to citizen science project?

12:00-1:40: Lunch: Huxley Building Room 217-218

Session 2

Huxley Building Room 144

1:40 - 2:00 Ilya Reshetouski
Kaleidoscopic Imaging

2:00 - 2:20 Grani Hanasusanto
Distributionally Robust Joint Chance Constraints with Conic Dispersion Measures

2:20 - 2:50 Chrysos Grigoris
Deformable face tracking in arbitrary videos

2:50 - 3:10 Thomas Whelan
ElasticFusion: Dense SLAM Without A Pose Graph

3:10 - 3:40 Coffee and voting for best presentation: Huxley Building Room 217-218

3:40 - 3:50 Prize giving for best presentation

Blackett Lecture Theatre 1

4:00 - 5:00 Keynote Speaker: Guido van Rossum *25 Years of Python*

Session 1

Speaker: Jeremy Cohen

Title: libhpc: A framework and tools for simplifying running of large-scale scientific HPC jobs for end-users

Abstract: Large-scale scientific computing is complex, both for the developers building software and for the end-users who need to run this software to support their research. As hardware continues to evolve, with multi-/many- core platforms, GPUs, and cloud computing infrastructure becoming widely used in HPC environments, taking advantage of the added capabilities that these platforms provide is a major challenge. From the perspective of end-user scientists, understanding how to run code on heterogeneous resources is often too complex or time- consuming. Even when using local cluster or server resources, these users may rely on support from domain developers, computer scientists and hardware platform operators or providers when running their jobs. Developers invest a significant amount of knowledge and understanding into the process of building scientific code and this is often lost within the code once an application is complete. We consider that maintaining this information as metadata is key in supporting long-term maintenance and sustainability of scientific codes. It can also be used effectively in providing abstractions that support the ease of use required by end-users while addressing the complexity of ensuring efficient execution across a range of computational environments.
The EPSRC-funded libhpc Stage II project is building the libhpc framework which is designed to provide a flexible e-Infrastructure to handle the running of user jobs on heterogenous hardware while decoupling and simplifying the complex interactions between scientists, developers and resource providers that are often required in such scenarios. Libhpc II, and its stage I predecessor, have developed core functionality that is used in a range of use cases to support different scientific domains. At the user level we have developed the concept of software parameter templates and job profiles that allow entities to collaborate on producing job configurations which can then be shared, edited and extended via a web-based interface. The libhpc II project is a collaboration between the departments of Computing (Darlington / Cohen) and Aeronautics (Sherwin), the Bioinformatics Support Service (Butcher) and the School of Public Health (Aanensen) at Imperial and Edinburgh Parallel Computing Centre, University of Edinburgh (Guo).

This presentation will provide a brief overview of libhpc and the techniques used in the framework and then demonstrate two environments for supporting users in running large-scale scientific applications - Nekkloud, an interface for running jobs using the Nektar++ spectral h/p element framework (http://www.nektar.info) and a pipeline builder for running molecular dynamics (MD) jobs using the GROMACS (http://www.gromacs.org) MD code. Both these tools make use of the template and profile model and demonstrate how this can be used to provide users with quicker and more effective means of configuring jobs.

Speaker: Poonam Yadav

Title: Interested in converting a scientific problem to citizen science project?

Abstract: Citizen Science projects are popular and easily accessible community science projects where volunteers can actively participate in science using web-based platforms. Alternatively, volunteers can participate passively by contributing computational power to projects. The process of setting up of a citizen science project requires careful considerations, and can be a challenging task for scientists. In this presentation, we present a systematic framework, which helps scientists to understand whether citizen science is a suitable approach to their research problem. It also helps them to understand where their application fits in the citizen science domain. This process guides them in setting up their own citizen science projects. We present a brief summary of available tools and techniques that make deployment of the project simple and easy process.

Speaker: Pavel Burovskiy

Title:

Abstract: will present the current progress in mapping Finite Element Methods (FEM) to the reconfigurable hardware (FPGAs). This effort aims to build a custom hardware accelerator to the HPC problem of a wide applicability. The main challenge addressed here is performing an unstructured gather/scatter on custom hardware more efficiently than CPU could.

Speaker: Lawrence Mitchell

Title: Symbolic numerical computing: solving equations by writing them down

Abstract: Huge swathes of scientific research are predicated on the ability to efficiently simulate physical phenomena, often modelled via systems of partial differential equations (PDEs). The ability to answer real-world questions necessitates the application of cutting edge numerical and computational techniques. Previous state of the art in such model development required the programmer to simultaneously be an expert in high performance computing, numerical analysis, along with their application domain.

In this talk I discuss a better way: the synthesis of high performance models for the numerical simulation of PDEs from symbolic problem descriptions. By carefully choosing abstractions appropriate to the mathematical problem domain, synthesis of low-level, high-performance code is possible: the domain expert need no longer be a polymath! This approach is concretely realised in the Firedrake project a python library for the automated solution of PDEs using the finite element method, and I will illustrate the approach with a number of examples.

Speaker: Luigi Nardi and Zeeshan Zia

Title: Multi-space multi-objective design-space exploration for 3D scene understanding using active learning

Abstract: Real-time dense computer vision offers great potential for a new level of 3D scene modeling, tracking and environmental interaction for many types of robots and augmented reality systems. However, their high computational requirements make actual applicability on embedded platforms challenging. While classical point feature-based simultaneous localisation and mapping (SLAM) are emerging in embedded systems, dense SLAM algorithms are still largely under investigation as research prototypes on desktop environments because their computational requirements prevent deploying them in the embedded space.

In this work we study a dense SLAM algorithm, namely KinectFusion, in the context of the SLAMBench framework and study how KinectFusion can be mapped to power constrained embedded systems. A key idea of our approach is the co-design space exploration of how algorithmic/compiler/architecture choices affect the power, runtime and accuracy of the underlying system. The total number of configurations of all the three spaces consists of millions of points that can not be explored exhaustively. We employ active learning techniques to prune the design-space allowing us to jointly explore the three spaces and find the optimal pareto front using a multi-objective machine learning optimiser.

Our results indicate that we can sustain a real-time frame rate (30 FPS) within a power budget of 1W on a popular embedded platform. This is a 6x improvement in performance compared with state-of-the-art systems. This outcome shows the suitability of dense SLAM systems for real-time embedded domains such as mobile robots.

Speaker: Nicholas Ng

Title: Pabble, Session Types and dataflow programming

Abstract: Pabble is a protocol description language for writing interaction-based communication protocols. It is an implementation of the formal typing discipline of Multiparty Session Types (MPST), which guarantees the absence of communication mismatch and deadlocks, and is a great way of abstracting structured communication patterns. In our previous work we used Pabble communication protocols to generate parallel programs by combining with sequential code snippets, which are specified as separate entities. With the emergence of large-scale data processing frameworks such as Apache Storm and Google's Cloud Dataflow, the dataflow model of computing is becoming increasingly relevant as a high-level abstraction of computation. In our experience with Pabble and its representation of communication, we find that they share some characteristics with the dataflow as computation; this have prompted us to explore the possibilities in interpreting Pabble for dataflow programming. I will talk about the challenges and the progress of our endeavour so far, and discuss the prospects of using Pabble or more generally, Session Types, as a unified model for describing both communication and computation.

Session 2

Speaker: Ilya Reshetouski

Title: Kaleidoscopic Imaging

Abstract: With a simple three-mirrors kaleidoscope it is possible to generate multiple views of the same scene in a single camera image. This property is very attractive in multi-view computer vision applications. However, for successful interpretation of such images it is important to distinguish different viewpoints, which may overlap in the recorded image. Image labeling is the assignment of viewpoints to image pixels. Successful labeling enables the extraction of single view point images from the recorded multiple-bounce mirror image.

In my talk I will present our solution of the labeling problem and demonstrate how labeling can be produced for cameras and projectors. Finally, I will show how we use labeled camera and projector images for geometry and reflectance acquisition.

Speaker: Grani Hanasusanto

Title: Distributionally Robust Joint Chance Constraints with Conic Dispersion Measures

Abstract: We analyse the complexity of a class of distributionally robust joint chance constrained programs where the uncertain parameters are described through their mean values and through upper bounds on general dispersion measures. We derive a tractable problem reformulation when the dispersion measure is conic and uncertain parameters only affect the right-hand side vector of the chance constraint. We also show that the problem becomes intractable if the left-hand side coefficient matrix is affected by uncertainty or the support of the uncertain parameters is restricted to a polyhedron. We illustrate the effectiveness of our exact reformulation in numerical experiments and demonstrate its superiority over state-of-the-art approximation schemes.

Speaker: Chrysos Grigoris

Title: Deformable face tracking in arbitrary videos

Abstract: Arguably, generic face detection and facial landmark localization in arbitrary static imagery are among the most mature and well-studied problems in the intersection of statistical machine learning and computer vision. Currently, the top performing face detectors achieve a true positive rate of about 75-80% when a small number of false positives is allowed. Furthermore, the top performing facial landmark localization algorithms can obtain a satisfactory low point-to-point error for more than 70% of a database's images that are captured under unconstrained conditions. The task of facial landmark tracking in videos has attracted much less attention. The pipeline that is commonly applied is no different than landmark detection in static imagery in order to avoid drifting. That is, a tracking-by-detection framework is applied where face detection and landmark localization are employed in every frame. Empirically, a straightforward application of such a framework could not achieve higher performance, on average, than the performance reported for static imagery. In our work, we show for the first time, to the best of our knowledge, that it is possible to use the results of generic face detection and landmark localization algorithms to recursively train powerful and accurate person-specific face detectors and landmark localization methods. We show that the proposed pipeline can track landmarks in very challenging long-term sequences captured under arbitrary conditions.

Speaker: Thomas Whelan

Title: ElasticFusion: Dense SLAM Without A Pose Graph

Abstract: We present a novel approach to real-time dense visual SLAM. Our system is capable of capturing comprehensive dense globally consistent surfel-based maps of room scale environments explored using an RGB-D camera in an incremental online fashion, without pose graph optimisation or any post-processing steps. As well as this, our system scales to larger scale "corridor-like" environments more traditionally seen in SLAM system experiments. In our evaluation we show that our approach improves both the trajectory estimate of the camera and the surface reconstruction quality, out performing many different existing approaches considered to be state of the art.

Keynote Speaker: **Guido van Rossum**

Title: 25 Years of Python

Abstract: This talk reviews Python's history, design philosophy, evolution and community, and gives a peek into future developments.

Short Bio: Guido van Rossum created the open-source programming language Python, and is its lead developer and thought leader ("BDFL"). Guido was born in the Netherlands and developed Python while at CWI in Amsterdam. He now lives in California where he works for Dropbox. He has previously worked for Google. He is an ACM Distinguished Engineer.