THE UNIVERSITY of York

Department of Computer Science

# Real-Time Systems Group

The Real-Time Systems Research Group at the University of York is now one of the largest academic research groups in the world focusing on the engineering of real-time systems.

## What is Real-Time?

A system is defined as being real-time if it is required to respond to input stimuli within a finite and specified time interval. The stimuli being either an event at the interface to the system or some internal clock tick that is, at least notionally, coordinated with the passage of time in the system's environment. Real-time systems are found in a wide range of applications areas, from simple domestic appliances to multimedia systems, large scale process control and safety critical avionics. In some systems the required response times are measured in milliseconds, in others it is seconds or even minutes. Nevertheless they all have deadlines that must be satisfied. In the production of real-time systems, it is insufficient to use testing of the final system to ensure compliance its with the timing requirements. A comprehensive and systematic approach to specification, design, implementation and analysis is required.

## Areas of Activity

The overall goal of the group is to facilitate the design, construction, analysis and maintenance of potentially complex systems which have real-time constraints. The work of the group spans a wide range of topics associated with timing analysis, system design, programming languages, operating system kernels, heterogeneous and reconfigurable hardware.

## Projects and Funding

Areas of application of our work include space and avionic systems, engine controllers, automobile control and multi-media systems. Work has been funded by the EPSRC and DTI, BAE SYSTEMS, European Union, European Space Agency (ESA), NASA, QinetiQ, Rolls Royce Aeroengines, DTI, the Health and Safety Executive (HSE), Sun Microsystems, Philips Research, Microsoft.

**Aim of the Group:** To undertake fundamental research, and to bring into engineering practice modern techniques, methods and tools.

How can we design and build time-predictable systems using modern heterogeneous multicore architectures?

By what means can a system utilise free resources at run-time, to increase system utility, but still be guaranteed to meet its fundamental timing constraints?

Can a computational model for real-time systems be defined that is effective in terms of schedulability analysis, but can also be formally specified and implemented?

How can we implement, in combinations of software and hardware, complex real-time systems to meet timing and resource constraints?

To what extent can real-time systems be made architecture neutral?

## Contact and Further Information

Alan Burns
Alan.Burns@cs.york.ac.uk
+44 1904 432779

Andy Wellings
Andy.Wellings@cs.york.ac.uk
+44 1904 432742

http://www.cs.york.ac.uk/rts

Department of Computer Science, University of York, Heslington, York - YO10 5DD, United Kingdom

## Permanent Academic Staff

**Neil Audsley**, Senior Lecturer: Implementation of predictable high performance embedded real-time systems: heterogeneous architectures and memory systems, operating systems, programming languages and compilers.

**Iain Bate**, Senior Lecturer: WCET, platform optimisation, co-design issues (HW/SW and control/scheduling), industrial applications.

**Alan Burns**, Professor: Scheduling theory, languages, architectures including kernels and communication protocols, dependability.

**Rob Davis,** Senior Research Fellow: Flexible real-time scheduling.

**Leandro Soares Indrusiak**, Lecturer: Model-driven design of embedded systems, on-chip multiprocessors, on-chip interconnect networks, wireless sensor networks.

**Andy Wellings**, Professor: Languages, Real-time Java, Ada, kernels, safety kernels, architecture-neutral systems, distributed programming.

## Research Staff (September 2010)

Emad Al-Oqayli, Research Student: Shared memory real-time non blocking communication, transactional memory.

Mohammed Al Rahmawy, Research Student: Real-time Java component models.

Mark Bartlett, Research Associate: Machine learning and WCET analysis.

Guillem Bernat, Visiting Research Fellow: Scheduling theory, worst-case execution time analysis, probabilistic techniques for timing analysis.

Andrew Burkimsher, Research Student: QoS-aware scheduling of HPC workflows.

Ipek Caliskanelli, Research Student: Service availability in sensor networks.

Mark Fairburn, Research Student: Knowledge centric design of wireless sensor networks.

Dave George, Research Student: Operating systems for heterogeneous architectures

Ian Gray, Research Student: Programming heterogeous systems, compilers

David Griffin, Research Student: Probabilistic WCET analysis.

Hashem Ghazzawi, Research Student: Control theoretic approach to scheduling HPC workflows.

Min Seong Kim, Research Associate: Real-time event handling, RTSJ.

Se-Yeon Kim, Research Student: NUMA operatiing systems, memory management.

Hui Keng Lau (Kelvin), Research Student: Real-time anomaly detection in complex systems.

Tiong Hoo Lim, Research Student: Lightweight fault tolerance in sensornets.

Gareth Lloyd, Research Student: Predictable Dynamic reconfiguration for real-time systems on FPGAs

Shiyao Lin (Charlie), Research Student: Multiprocessor shared memory communication

Abdul Haseeb Malik, Research Student: Real-time Java, NUMA architectures.

Gary Plumbridge, Research Student: High level language synthesis for FPGA accelerators

Tom Richardson, Research Student: Real-time component engineering.

Kun Wei, Research Assistant: Time bands for real-time systems.

Sitsofe Wheeler, Research Student: Machine learning, WCET analysis.

Jack Whitham, Research Associate: Predictable architectures and memory systems.
Alex Wood, Part-Time Research Student: WCET and machine learning.

Ke Yu, Research Student: Simulation of embedded systems.

Attila Zabos, Part-Time Research Student: Mode-change protocols, flexible scheduling.

Scheduling Theories

Design Synthesis

Static Code Analysis

Worst-case Execution Time (WCET) Analysis

Search-Based Engineering

Distributed & Parallel Architectures

Concurrent Object-Oriented Languages

Design Methods

Probabilistic Theories of Time

Network on a Chip

Sensor Networks

Real-Time Java

Fault Tolerance

Technology Transfer

Kernels

Embedded Systems

Predictable Heterogeneous Architectures and Memory Systems

Portable Code

Ada