

BT Research contacts: Fabrice Saffre & Mark Shackleton

Problem statement: A key issue in Information Communication Technology (ICT) is that of deciding where a distributed set of processes should run. This is driven by the fact that ICT is increasingly distributed in nature with the advent of "Grid" computing and "Cloud" computing models. *The problem is really one of optimising the geographic distribution of many individual processes that need to execute against a highly dynamic and decentralised set of available resources.* The optimisation criteria include "quality of service" (QoS) factors such as response time as well as other dynamically changing factors such as local costs and energy factors. These factors lead to interesting tradeoffs and conflicting demands that need to be taken into account during optimisation.

Nature inspired optimisation techniques seem particularly relevant to such a complex problem, and we propose two related projects, each of which investigates a different nature-inspired optimisation technique. The **mini-project** stage will focus on formalising the problem statement, defining the chosen methodology, and carrying out a literature survey of related approaches that have been used on such problems. The **summer project** will build on this initial research to create software that embodies the chosen nature-inspired optimisation technique and explores the parameter space when optimisation factors are weighted differently and the effect that has on the solutions attained.

As *we propose two sister projects*, there would be scope for two students to cross-compare their alternative solutions, and to collaborate on some aspects of the problem, such as problem and objective functions definitions and software.

Project 1: This project will focus on the use of **Swarm computing** to allocate processes (services) across the range of dynamically available resources. This will build on some research already carried out at BT and investigate how the migration of processes can be coordinated in an adaptive and distributed manner to match dynamically changing resource availability.

Project 2: This project will focus on the use of **evolutionary computing**. Different mappings between processes and resources will be encoded as genomes, on which the usual operators (selection, mutation, cross-over) can be applied. The selection operator will take into account both local factors (likely to be dominated by the cost of hosting resources) and global factors (like the cost of interaction between distant service components). As an additional subtlety, the fitness landscape will vary over time, as interaction patterns and resource availability fluctuate. The potential use of runtime adaptation in this dynamic context will be investigated.

Skills: students will need to draw upon the skills learnt during their Msc but will also need to have good Java coding skills to realise the required software.