MathSys / Complexity mini-project:
Simulation Optimisation with Simulated Annealing

Background: Simulation has become an invaluable tool for analysing and understanding complex systems. The natural next step is to use simulation not only to analyse a complex system, but to optimise it, by automatically searching through a vast array of possible alternative configurations. This is known as "simulation optimisation", and often black box optimisation heuristics such as evolutionary algorithms or simulated annealing are used for the search.

However, if the simulation model is stochastic and multiple replications lead to different results, optimisation becomes difficult, as it is not clear whether an observed quality difference is real, or only due to randomness in evaluation. Averaging over multiple replications reduces the effect of noise, but quickly becomes computationally intractable. Thus, there is a need for efficient optimisation algorithms that only require a minimal number of simulation replications – enough to ensure progress in optimisation, but not more than necessary in order to limit the required computational time. Initial attempts to design such algorithms can be found for example in [1,2]. Recently, we have developed another variant of simulated annealing that, under certain assumptions, only takes the provably minimal number of simulation replications (paper under review).

Mini-project: In this project, we collaborate with Lanner Ltd, the company behind WITNESS a leading simulation software (http://www.lanner.co.uk/en/witness.cfm). Lanner currently sells its own version of simulated annealing for simulation optimisation. One goal of the miniproject would be to compare the performance of our new algorithm with Lanner's current algorithm. Another goal would be to further relax some of the assumptions currently made and thus make our new algorithm applicable to a wider range of problems.

PhD prospect: The topic offers ample opportunities for extension into a PhD thesis, including making it more practical to real-world scenarios (e.g., currently we assume the standard deviation of the noise is known, but in practice it has to be estimated during the optimisation run), extending it to other optimisation criteria (e.g., currently we optimise the expected value, but optimising a quantile may be interesting), or parallelising the algorithm so that it can run in the cloud.

Deliverables: Algorithm comparison and extension.

Student's requirements: Some programming is necessary.

Supervisors: Prof Juergen Branke, Prof Robin Ball

References

[1] Ball, R. C., T. M. A. Fink, N. E. Bowler. 2003b. Stochastic annealing. Physical Review Letters 91(3) 030201

[2] Branke, J., S. Meisel, C. Schmidt. 2008. Simulated annealing in the presence of noise. Journal of Heuristics 14:627–654