

Optimization in eScience with Nimrod/O

Tom Peachey and David Abramson

Faculty of Information Technology,
Monash University, Australia.

A considerable part of eScience involves the repeated execution of computational models in order to determine that combination of input parameters that will optimize some aspect of the output. One such set of problems involve industrial design. Increasingly, in the design of engineering machines and structures, the prototyping stage is being replaced by computer modelling. Typically each design evaluation requires models of the continuum such as finite element models, computational fluid dynamics, models of diffusion, or electromagnetic fields. A design may be required that meets the design criteria and maximizes the performance in some sense or minimizes the cost.

A second class of problems occurs in scientific modelling where a computer model simulates a physical system. These often present the user with an inverse problem, finding the input parameters that produce an output like that of the real system modelled. This can be tackled as an optimization problem, minimizing the discrepancy between the model output and real world data.

Typically the models are computationally intensive and an optimization requires many evaluations of the model, so the computational task is substantial. Nimrod/O is a tool that automates such tasks. It supplies a range of standard optimization search algorithms and expedites the solution by evaluating jobs in concurrent batches. A guiding principle in the design of Nimrod/O is that the user should not need knowledge of distributed computing nor of the details of optimization methods. Nimrod/O offers an "off the shelf" selection of methods with sensible default settings. For a given class of optimization problems it is unknown in advance which algorithms will be best suited; Nimrod/O allows the user to perform a variety of search algorithms concurrently. Further, some problems are sensitive to the starting point chosen for the search. Nimrod/O can perform concurrent multiple starts, limited only by the number of processors available.

Often the researcher has already developed an effective optimization method but wishes to avail of the parallelism of Nimrod/O. A recent addition is a library of functions that allows external programs to be co-scheduled and to access the caching and concurrency of Nimrod/O.

Nimrod/O is being or has been used in a wide range of eScience projects.

- Design of mechanical components for optimal fatigue life.
- GAMESS - optimization of the effective group difference pseudopotential in a quantum chemistry model.
- Determination of the operator norm for the generalized Hilbert transform.
- Models of star formation

- Modelling geophysical extension models
- Design of an airfoil to optimize the ratio of lift to drag.
- Electromagnetic modelling for design of an optimal test rig for a mobile telephone antenna.
- Air quality modelling with the aim of minimizing ozone pollution.