Enhancing the Performance Predictability of Grid Applications with Patterns and Process Algebras

The ENHANCE Project

Anne Benoit, Murray Cole, Stephen Gilmore, Jane Hillston and Gagarine Yaikhom

http://groups.inf.ed.ac.uk/enhance/

Overview

One of the most promising technical innovations in present-day computing is the invention of Grid technologies which harness the computational and storage power of widely distributed collections of computers. Grid technologies are breaking new ground in e-Science where scientists across the globe can collaborate and share data sets, processing power and specialised scientific instruments to accelerate the pace of scientific development. Grid technologies are also opening new doors in e-Business where small or medium-scale businesses can now have access to supercomputing power which was once the province of only the most wealthy banks and multi-nationals. Data and resource sharing schemes such as these are changing the basic ground rules of computing.

Such systems raise difficult issues of resource allocation and scheduling (roughly speaking, how to decide which computer does what, and when, and how they interact) which are made all the more complex by the inherent unpredictability of resource availability and performance. For example, I may forget to switch on my PC, a supercomputer may be required for a more important task, the internet connections I need may be particularly busy.

The Enhance project aims to simplify the effective programming of Grid systems by exploiting and synthesising results from two underlying research programmes. Stochastic process algebras such as PEPA are used to model the behaviour of concurrent systems in which some aspects of behaviour are not precisely predictable. Meanwhile, pattern based programming recognises that many real applications draw from a range of well known solution paradigms and seeks to make it easy for an application developer to tailor such a paradigm to a specific problem without re-inventing the wheel. Our key insights are:

- that the use of patterns will helpfully constrain the implementation challenge, by providing good knowledge of the overall evolution of an application's interactions;
- correspondingly, that patterns allow both more accurate cost prediction (which in turn informs good scheduling decisions) and lighter weight monitoring (because the pattern based model tells us what the crucial aspects are for each instantiation);
- that cost modelling with a formalised, stochastic performance process algebra and associated tools facilitates capture and manipulation of the remaining uncertain aspects of application behaviour.

Case Study

We have undertaken several case studies to check the practical usefulness of our methods. One of these relates to the on-line scheduling of a numerically-intensive algorithm which uses the Mean Value Analysis approach to compute metrics of a queueing network. The algorithm is a pipeline-structured parallel application implemented using C, MPI and the Edinburgh skeletons library.

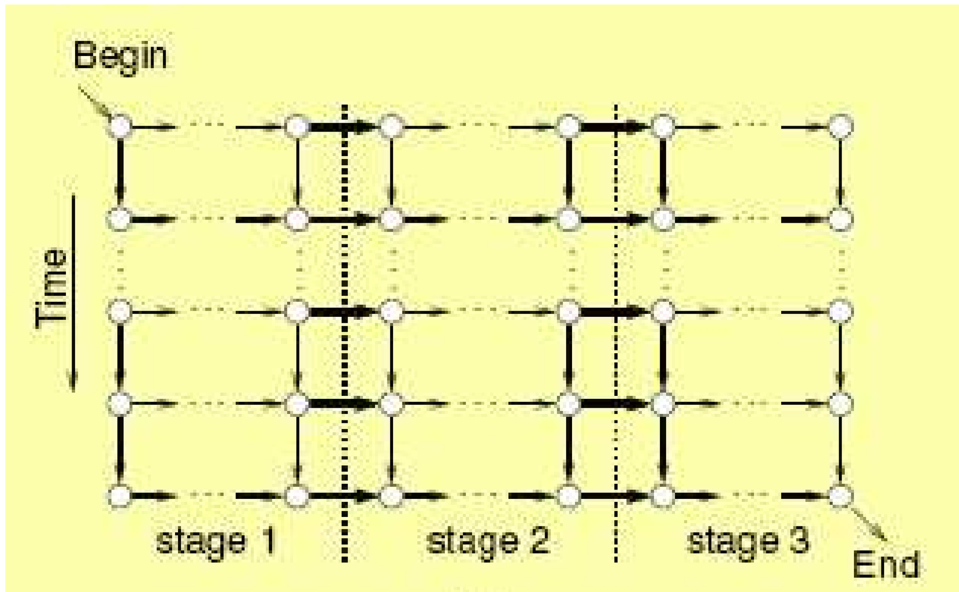

Our system dynamically samples the CPU availability and inter-node TCP communication latency of our Beowulf compute cluster. It then uses the measurements so obtained to parameterise an automatically generated PEPA model of the pipeline-structured parallel application.

$$
\begin{aligned}
Source_s &\stackrel{def}{=} (move_s, \top).(input_{s,1}, \top).\\
&\quad (move_s, \top).(input_{s,2}, \top).\\
&\quad \ldots\\
&\quad (move_s, \top).(input_{s,N_{w_s}}, \top).Source_s\\
Worker_{s,i} &\stackrel{def}{=} (input_{s,i}, \top).(process_{s,i}, \top).(output_{s,i}, \top).Worker_{s,i}\\
Sink_s &\stackrel{def}{=} (output_{s,1}, \top).(move_{s+1}, \top).\\
&\quad (output_{s,2}, \top).(move_{s+1}, \top).\\
&\quad \ldots\\
&\quad (output_{s,N_{w_s}}, \top).(move_{s+1}, \top).Sink_s\\
Stage_s &\stackrel{def}{=} Source_s \underset{LI_s}{\bowtie} (Worker_{s,1} \parallel \ldots \parallel Worker_{s,N_{w_s}}) \underset{LO_s}{\bowtie} Sink_s
\end{aligned}
$$

The system solves the PEPA model using the PEPA Workbench, and derives performance metrics which predict which scheduling decisions will lead to the shortest run-time. The MVA application is then scheduled accordingly.

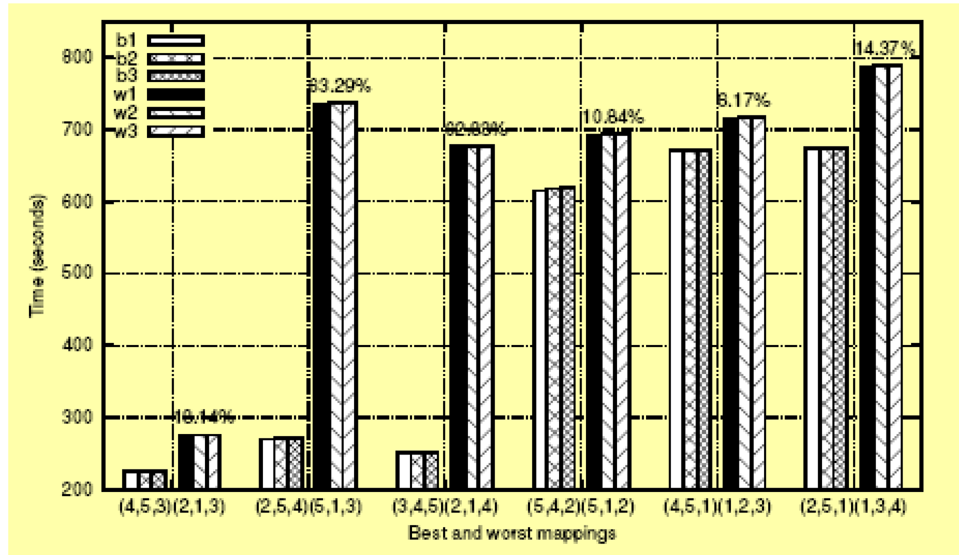

In the present study we did not find a single instance where the performance predictions computed by the PEPA Workbench were misleading: predicted best cases always significantly outperformed predicted worst cases in practice. Strategic use of these methods will reduce the run-times of these parallel codes overall, making better use of heavily-loaded compute clusters.

In the latest version of our system, we are able to dynamically adjust the schedule in response to varying external load conditions on the processors and network.

Sample Publications

Anne Benoit, Murray Cole, Stephen Gilmore and Jane Hillston, "Evaluating the performance of pipeline-structured parallel programs with skeletons and process algebra", *Scalable Computing: Practice and Experience*, vol 6, no 4, pp 1–16, 2005.

Gagarine Yaikhom, Murray Cole, and Stephen Gilmore, "Combining measurement and stochastic modelling to enhance scheduling decisions for a parallel Mean Value Analysis algorithm,", *Proceedings of 6th International Conference on Computational Science (ICCS 2006)* Springer-Verlag LNCS 3992, pp 929–936, 2006.

Funding

The Enhance project is funded by the Engineering and Physical Sciences Research council grant number GR/S21717/01.

School of informatics