Performance Prediction of Cloud Applications Using Dwarf Benchmarks

Thursday, 20 September 2012 11:20 (20 minutes)

Description of the work

We have developed a benchmark suite for measuring the performance of virtualised hardware based on the earlier work carried out in the TORCH project [1]. This collection uses the concept of algorithmic patterns, in the literature referred to as "dwarfs". "A dwarf is an algorithmic method that captures a pattern of computation and communication"[2]. The proposed dwarf taxonomy [2] consists of 13 frequently occurring patterns that can be considered as building blocks of applications, such as dense matrix operations, map reduce, particle and spectral computations. This feature makes them especially suitable for performance modelling. For validating the prediction technique we have included applications from different domains such as molecular dynamics (Gromacs [3]), rendering (Blender [4]) and video transcoding (FFMPEG [5]). We have carried out extensive benchmarking on BonFIRE [6] project resources and on facilities of commercial Cloud providers such as Amazon EC2, Rackspace, CloudSigma, GoGrid and ElasticHosts. As a mathematical apparatus for modelling initially we used polynomials and later the Moore-Penrose inverse matrix [7] that enabled us to construct multi-dimensional models by combining the data from all dwarfs and has significantly improved the accuracy of predictions. Our work has demonstrated that the concept of algorithmic dwarfs can be successfully used not only for the characterisation of Cloud resources but also for performance modelling. We have also investigated the problem of reliability of long term experiments and the impact of various levels of contention on processes running on virtual resources.

Link for further information

http://eprints.soton.ac.uk/273157/

Wider impact of this work

The concept of algorithmic dwarfs has been presented a few years ago. However, the modelling power of this approach has not been put to a rigorous test until now. We have collected data from several cloud providers and validated this concept on four applications drawn from different computational domains. We believe that our work can be used by several stakeholders in Cloud market for:

a) making better provisioning decisions: deploying the infrastructure resources required for a given application QoS rather than over-provisioning.

b) making better application scheduling decisions: knowing the application runtime with a good reliability permits more intelligent scheduling.

c) determining the optimal application configuration: the performance of complex applications and business or industrial data processing workflows with many components can be greatly affected by their configuration.

Printable Summary

The EC-FP7 BonFIRE project offers a multi-site Cloud facility of geographically distributed testbeds of heterogeneous resources for Future Internet experimentation. IT Innovation has taken part in the project as one of three driving experiments, focusing on QoS-oriented service engineering for federated Clouds. In this presentation, we will discuss the challenges we have addressed pertaining to resource provisioning for applications that have specific QoS requirements and where consumers of Cloud resources want to avoid lock-in to any specific IaaS provider. This is a significant challenge due to the limited and varying information IaaS providers disclose about infrastructure resources. This presentation covers the use of application-focused benchmarks, referred to as Dwarfs, for characterising compute resources and for application performance modelling. This work benefits both researchers and software developers involved in the SaaS, Paas and IaaS areas of Cloud computing.

Primary author: Mr BONIFACE, Michael J. (IT Innovation Centre, University of Southampton)

Co-authors: Mr PAPAY, Juri (IT Innovation Centre, University of Southampton); Mr PHILLIPS, Stephen C. (IT Innovation Centre, University of Southampton); Mr ENGEN, Vegard (IT Innovation Centre, University of Southampton)

Presenter: Mr PAPAY, Juri (IT Innovation Centre, University of Southampton)

Session Classification: Virtualised Resources

Track Classification: Virtualised Resources: challenges and opportunities (Michel Drescher: track leader)