Energy conservation in distributed heterogeneous computing environments using economic resource allocation mechanisms

by

Timothy Michael Lynar

BIT(Hons)

A thesis submitted in partial fulfilment of the requirements for the degree of

Doctor of Philosophy

The University of Newcastle

June, 2011

This thesis contains no material which has been accepted for the award of any other degree or diploma in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. I give consent to this copy of my thesis, when deposited in the University Library, being made available for loan and photocopying subject to the provisions of the Copyright Act 1968.

Signature______________________________

Date______________________________

Abstract

Energy conservation in distributed heterogeneous computing environments using economic resource allocation mechanisms

by Timothy Michael Lynar

This thesis examines the question: can economic resource allocation mechanisms be used in distributed computing environments to reduce energy consumption whilst maintaining execution speed? This thesis investigates the use of several resource allocation mechanisms that take account of the power consumption and processing capacity of each available computing node within a distributed heterogeneous computing environment. Different economic resource allocation mechanisms have different attributes and allocate resources differently. The resource allocation mechanisms are evaluated to examine their effect on the time and energy required to process a workload of the sort that might be expected in a distributed computing system. Initial examination of the resource allocation mechanisms was conducted through the execution of artificial workloads on a simulated cluster. To further this research, a real cluster and grid environment was created from obsolete computers. An examination was undertaken of the use of obsolete computers in distributed computing environments and how the use of such systems may assist to mitigate electronic waste. The examination of resource allocation was continued on a cluster, and then on an institutional grid. The simulation model was then calibrated to the cluster and grid, which was then used to simulate the execution of real published grid workloads under each of the resource allocation mechanisms. The resource allocation mechanisms under consideration were found to have different characteristics that resulted in them being suited for different types of workload. It was also found that the choice of a resource allocation mechanism that takes account of the power consumption and performance of individual resources can make a significant difference, through leveraging the heterogeneous nature of resources, to the total system energy consumed and time taken in computing a workload.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

Thank you to my supervisors Dr. Ric Herbert and Dr. William Chivers, for their help and encouragement; to the Blue Gum Flats research group for their encouragement; to Simon for the help and direction he provided, and for assistance with English expression. Thank you to Steve Owers and Gary Maynard for their assistance in acquiring space and equipment. Thank you to Nic Croce and Maureen Townley-Jones from the Statistical Support Service from the Faculty of Science and Information Technology for their time and statistical advice. My sincere thanks to Rod Bell for providing invaluable feedback on this thesis. I acknowledge the help of a professional thesis editor Belinda Leskiw.

DEDICATION

To my wife Rachel,

who provided the most encouragement and support of all.

LIST OF ABBREVIATIONS

BA: Batch auction

CDA: Continuous double auction

CLUSTER: Computing cluster

CRA: Continuous random allocation

DVFS: Dynamic voltage and frequency scaling

EOE: End of execution

E-WASTE: Electronic waste

FLOPS: Floating point operations per second

FOSS: Free and open source software

GT: Grid trace

HPC: High-performance computing

HPDC: High-performance distributed computing

HPL: High-performance Linpack benchmark

MFLOPS: Millions of floating point operations per second

MIM: Marginal increase mechanism

MPI: Message passing interface

MWIPS: Millions of whetstone instructions per second

NODE: Computing node

PPBA: Pre-processed batch auction

PXE: Pre-execution environment

SPMD: Single process, multiple data

TTT: Test task time

UML: Unified modelling language

VFS: Voltage and frequency scaling

VOVO: Variable-on variable-off

WF: Workflow

WIPS: Whetstone instructions per second

ZIT: Zero intelligence trader

PUBLICATIONS RESULTING FROM THIS RESEARCH

Fully refereed journal articles

Lynar, T., Simon, Herbert, R. D., & Chivers, W. J. (2011). Resource allocation to conserve energy in distributed computing. International Journal of Grid and Utility Computing (IJGUC), 2 (1), 1-10.

Lynar, T. M., Herbert, R. D., Simon, & Chivers, W. J. (2010). Clustering obsolete computers to reduce e-waste. International Journal of Information Systems and Social Change, 1(1), 1-10.

Lynar, T. M., Herbert, R. D., & Simon. (2009). Auction resource allocation mechanisms in grids of heterogeneous computers, WSEAS Transactions on Computers, 10, 1671-1680.

Fully refereed conference proceedings

Lynar, T., Simon, Herbert, R. D., & Chivers, W. J. (2010). Reducing grid energy consumption through choice of resource allocation method. The 6th Workshop on High-Performance, Power-Aware Computing (HPPAC). Atlanta, GA, USA. April 19.

Lynar, T. M., Herbert, R. D., Simon, & Chivers, W. J. (2009). Impact of node ranking on outcomes of grid resource allocation. The 2009 International Conference on Grid Computing and Applications (GCA09). Las Vegas, NV, USA. July 13-16.

Lynar, T. M., Herbert, R. D., Simon, & Chivers, W. J. (2009). A grid resource allocation mechanism for heterogeneous e-waste computers. The 7th Australasian Symposium on grid computing and e-research (Ausgrid 2009). Wellington, New Zealand. January 20–23.

Herbert, R. D., & Lynar, T. M. (2009). A comparison of economic resource allocation mechanisms in grids of e-waste computers. The proceedings of the 9th WSEAS International Conference on Simulation, Modeling and Optimization (SMO'09). Budapest, Hungary. September 3–5.

Lynar, T. M., & Herbert, R. D. (2009). Allocating grid resources for speed and energy conservation. The 6th International Conference on Information Technology and Applications (ICITA09). Hanoi University of Technology. Hanoi, Vietnam. 9–12 November.

Lynar, T. M., Herbert, R. D., Simon, & Chivers, W. J. (2009). Why decide: Is a user's estimation of job completion time useful in grid resource allocation? The 18th World IMACS Congress and International Congress on Modelling and Simulation (MODSIM09). Cairns, QLD, Australia. 13–17 July.