

Characterizing and Evaluating Desktop Grids: An Empirical Study

Derrick Kondo¹ Michela Taufer^{1,3} Charles L. Brooks III³ Henri Casanova^{1,2} Andrew A. Chien¹

¹Computer Science and Engineering Department, University of California at San Diego

²San Diego Supercomputer Center

³Molecular Biology Department, The Scripps Research Institute

[dkondo, taufer, casanova, achien]@cs.ucsd.edu, brooks@scripps.edu

Abstract

Desktop resources are attractive for running compute-intensive distributed applications. Several systems that aggregate these resources in desktop grids have been developed. While these systems have been successfully used for many high throughput applications there has been little insight into the detailed temporal structure of CPU availability of desktop grid resources. Yet, this structure is critical to characterize the utility of desktop grid platforms for both task parallel and even data parallel applications. We address the following questions: (i) What are the temporal characteristics of desktop CPU availability in an enterprise setting? (ii) How do these characteristics affect the utility of desktop grids? (iii) Based on these characteristics, can we construct a model of server “equivalents” for the desktop grids, which can be used to predict application performance? We present measurements of an enterprise desktop grid with over 220 hosts running the Entropia commercial desktop grid software. We utilize these measurements to characterize CPU availability and develop a performance model for desktop grid applications for various task granularities, showing that there is an optimal task size. We then use a cluster equivalence metric to quantify the utility of the desktop grid relative to that of a dedicated cluster.