

XtremLab: A System for Characterizing Internet Desktop Grids

Paul Malecot Derrick Kondo Gilles Fedak
Laboratoire de Recherche en Informatique/INRIA Futurs
Bat 490, Université Paris Sud, 91405 ORSAY Cedex, FRANCE
[malecot, dkondo, fedak]@lri.fr

Abstract

Desktop grid (DG) systems use the idle computing power of many volunteered desktop PC's on the Internet to support large-scale computation and storage. For over a decade, DG systems have been the largest and the most powerful distributed computing systems, offering a plethora of computing power at a fraction of the cost of supercomputers. The volunteer desktops participating in DG projects are volatile and heterogeneous, but there is little detailed information about their volatility and heterogeneity. Yet this characterization is essential for the simulation and modelling of such systems. We are conducting a project whose goal is to obtain a detailed picture of the DG landscape. To this end, we are deploying resource sensors on the volunteer desktops to measure the availability of CPU, disk space, memory, and network bandwidth. An initial version of the project has been deployed (see <http://xtremlab.lri.fr>), and currently has about 1,800 participating hosts. The resulting resource measurement data and characterization will be useful for a broad range of research areas, including distributed and peer-to-peer computing, and fault tolerance.

I. Introduction

Since the late 1990's, desktop grid (DG) systems, such as SETI@Home, have been the largest and most powerful distributed computing systems in the world, offering an abundance of computing power at a fraction of the cost of dedicated, custom-built supercomputers. Many applications from a wide range of scientific domains – including computational biology, climate prediction, particle physics, and astronomy – have utilized the computing power offered by DG systems. DG systems have allowed these applications to execute at a huge scale, often resulting in major scientific discoveries that would otherwise had not been possible.

The computing resources that power DG systems are shared with the owners of the machines. Because the

resources are volunteered, utmost care is taken to ensure that the DG tasks do not obstruct the activities of each machine's owner; a DG task is suspended or terminated whenever the machine is in use by another person. As a result, DG resources are volatile in the sense that any number of factors can cause the task of a DG application to not complete. These factors include mouse or keyboard activity, the execution of other user applications, machine reboots, or hardware failures. Moreover, DG resources are heterogeneous in the sense that they differ in operating systems, CPU speeds, network bandwidth, memory and disk sizes. Consequently, the design of systems and applications that utilize these system is challenging.

II. Challenges and Opportunities

While DG systems consist of volatile and heterogeneous computing resources, it unknown exactly how volatile and heterogeneous these computing resources are. While there have been previous characterization studies on Internet-wide computing resources [2], these studies do not take into account causes of volatility such as mouse and keyboard activity, other user applications, and machine reboots. Moreover, these studies often only report coarse aggregate statistics, such as the mean time to failure of resources. Yet, detailed resource characterization is essential for determining the utility of DG systems for various types of applications. Also this characterization is a prerequisite for the simulation and modelling of DG systems in a research area where many results are obtained via simulation, which allow for controlled and repeatable experimentation.

III. Approach

We are currently designing, implementing, and deploying a resource monitoring project called XtremLab via the Berkeley Open Infrastructure for Network Computing (BOINC) [4] software system. BOINC projects are arguably the largest distributing computing projects in the world, as the number of active participants is more than 500,000 desktop PC's. The goal of XtremLab will be to

monitor the availability of a large fraction of these in an effort to paint a more detailed picture of the Internet computing landscape.

Specifically, the XtremLab application will periodically measure the availability of the CPU, memory, disk space, and network bandwidth, and write these measurements to a trace file. These trace files will then be collected and assembled to yield a continuous time series of availability for each hardware resource. We will then analyze these trace files to determine aggregate statistics of the system as a whole, and per-host statistics.

An initial version of the project has been deployed (see <http://xtremlab.lri.fr>), and currently (as of April, 2006) has about 1,800 participating hosts, from which we are currently collecting CPU availability trace data. To determine CPU availability, we use an active measurement method by which we submit real tasks to the system. Each task continuously computes integer and floating-point operations, and periodically writes the number operations to file. The advantage of this method is that the tasks perceive CPU availability exactly as a real DG application would, and this method is not prone to OS idiosyncrasies (by contrast to passive measurement methods [8], [3]).

IV. Impact

The results of this research will be useful to distributed computing research and other fields in many of ways. First, the trace data will enable accurate simulation and modelling of DG systems. For example, the traces could be used either to directly drive simulation experiments or to create generative probability models of resource availability, which in turn can be used by simulators to explore a wide range of hypothetical scenarios.

Second, because the traces will contain the temporal structure of availability, the traces will enable the assessment of the utility of DG systems for a wide range of applications. Currently, the range of applications that utilize DG systems effectively has been limited to applications with loosely-coupled tasks that are independent of one another; the volatility and heterogeneity of DG resources makes the execution of more tightly-coupled applications extremely challenging.

In addition, we believe our measurements could be useful for other sub-domains in computer science such as fault tolerance, peer-to-peer computing, and Grid computing. For example, one issue relevant to the fault tolerance research community is how often resources crash and why. The data we collect will reflect the time to failure for each desktop resource and thus be a valuable data set for those researchers. We will make the traces publicly available to all these research communities, such as those participating

in the French Grid5000 project and CoreGrid, a research organization for fostering European collaboration.

V. Past Work

We have conducted a number of related research efforts. First, in [7], [6], we measured and characterized several DG systems at the University of California at San Diego and the University of Paris-Sud. We obtained several months of traces of the availability of hundreds of desktop PC's within these organizations. We then characterized the DG systems by obtaining several aggregate and per-host statistics. This characterization formed the basis for a model describing the utility of the DG systems for different applications, and for developing efficient ways of scheduling tasks to DG resources [5]. So that others could use our gathered trace data sets, we created an online DG trace archive publicly accessible at <http://vs25.lri.fr:4320/dg>. The limitation of this work, which we address in XtremLab, is that no measurements were taken of home desktop PC's, which contribute significantly to Internet-wide DG projects.

Second, in [1], we measured aggregate statistics gathered through BOINC. A limitation of this work is that the measurements do not describe the temporal structure of availability of individual resources, and so it would be difficult to use these statistics for detailed simulation or modelling purposes. Again, our proposed work addresses this limitation by obtaining a time-series of measurements of DG resource availability.

References

- [1] D. Anderson and G. Fedak. The Computational and Storage Potential of Volunteer Computing. Technical report, Space Sciences Laboratory, UC Berkeley, November 2005.
- [2] R. Bhagwan, S. Savage, and G. Voelker. Understanding Availability. In *In Proceedings of IPTPS'03*, 2003.
- [3] P. Dinda. The Statistical Properties of Host Load. *Scientific Programming*, 7(3-4), 1999.
- [4] The Berkeley Open Infrastructure for Network Computing. <http://boinc.berkeley.edu/>.
- [5] D. Kondo, A. Chien, and Casanova H. Rapid Application Turnaround on Enterprise Desktop Grids. In *ACM Conference on High Performance Computing and Networking, SC2004*, November 2004.
- [6] D. Kondo, G. Fedak, F. Cappello, A. A. Chien, and H. Casanova. Resource Availability in Enterprise Desktop Grids. Technical Report 00000994, Dept. of Computer Science, INRIA, January 2006, submitted.
- [7] D. Kondo, M. Taufer, C. Brooks, H. Casanova, and A. Chien. Characterizing and Evaluating Desktop Grids: An Empirical Study. In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS'04)*, April 2004.
- [8] R. Wolski, N. Spring, and J. Hayes. The Network Weather Service: A Distributed Resource Performance Forecasting Service for Metacomputing. *Journal of Future Generation Computing Systems*, 15(5-6):757-768, 1999.