

EGEE User Forum

Contribution ID : 18

Fault Detection and Diagnosis from the Logging and Bookkeeping Data

Friday 11 May 2007 at 11:20 (00h20')

Description of the activity :

Autonomic Computing (AC) is defined as "computing systems that manage themselves in accordance with high-level objectives from humans" [1]. AC is now a well-established scientific domain, and a priority for industry. Automated detection, diagnosis, and ultimately management, of software/hardware problems define autonomic dependability. The paper reports on applying state of the art autonomic dependability methods to the Logging and Bookkeeping data, with promising results on detection.

Grid added value :

The production status and integration level reached by the EGEE middleware and monitoring provide immense datasets. These are challenging targets for the Machine Learning (ML) community, whose techniques are at the base of AC. The fundamental motivation for this interest is the complexity of the hardware/software components, and the intricacy of their interactions, which defeat attempts to build models only from a-priori knowledge. Furthermore, EGEE is not a steady-state system, not only because it is yet ramping-up, but more profoundly because of the externally-driven collective behaviour of its users.

EGEE monitoring data exemplify to the extreme two classical issues in ML: 1) curse of dimensionality (state space exponential in the number of variables); and 2) data sparsity, most of the state-action space being actually unexplored. EGEE data offer an extra complexity, not addressed in this paper: integration of heterogeneous sources of information (ontology building).

Experience or proposed activity :

L records are oriented towards operational semantics: each service logs its own vision of the job information and status. The result is a very large amount of highly redundant data, with in many cases no a-priori syntax or semantics (blobs in the long_fields table). We have first developed a software suite that segments the data in order to discover the basic attributes, and cautiously filter out redundant information. The software also allows to convert the categorical data into a boolean description, convenient for many off-the-shelf mining and

learning software. The next step was analysis. Elementary methods (independently scoring attributes) provided little information. The ROGER algorithm [2] developed in our lab provides a good predictor, which can be interpreted through sensitivity analysis. On-going work deals with intelligent clustering in order to reduce dimensionality (frequent itemsets) and learning non-linear models, which can detect compound failure conditions.

Future evolution :

This paper is an attempt to explore the performance and limits of this purely passive approach of detection and diagnosis: a minimally invasive failure analyser would be based solely on analysis of the production (L) records of the production jobs. A more intrusive approach would use SAM (or equivalent software) and active learning methods to (approximately) design an optimal probe set and infer the system state [3]. Support for gathering and interpreting data in this area would be required.

Primary authors : Prof. GERMAIN-RENAUD, Cecile (LRI and LAL (CNRS - University Paris-Sud 11)) ;
Ms. ZHANG, Xiangliang (LRI) ; Dr. SEBAG, Michèle (LRI (CNRS) and INRIA) ; Dr.
LOOMIS, Charles (LAL)

Co-authors :

Presenter : Prof. GERMAIN-RENAUD, Cecile (LRI and LAL (CNRS - University Paris-Sud 11))

Session classification : Grid Monitoring and Accounting

Track classification : Grid Monitoring and Accounting

Type : oral presentation