Software Fault Diagnosis for Grid Middleware with Bayesian Networks*

Jan Ploski
OFFIS Institute for Information Technology
Jan.Ploski@offis.de

Wilhelm Hasselbring
Carl von Ossietzky University of Oldenburg
Hasselbring@informatik.uni-oldenburg.de

Software failures after deployment consist of producing incorrect outputs or refusing to provide service altogether. In order to restore the expected service, people responsible for a failed application at a users' organization often have to infer an observed error's cause and the possible repair actions based on incomplete or even misleading information produced by the diagnosed software [BKM+04, RSB03].

Insufficient attention given to error handling during development is both easy to blame and to dismiss as the reason for the poor quality of error messages, being a project-specific human factor. However, more universal reasons for the poor quality of error messages exist that coincide with fundamental principles of software modularity. Specifically, the focus on software reuse and information hiding may lead to module interfaces with underspecified implementation-specific exceptions [PH05]. In general, module implementors may receive too little information from the execution environment to provide meaningful error messages.

In light of these issues, we propose an approach which supports fault diagnosis with a Bayesian network [Pea88]. The Grid middleware Condor [TTL05] served as an initial case study to test this approach within the e-Science project WISENT [WIS06]. The Bayesian network is constructed manually from a user's perspective in order to link each fault hypothesis to symptoms observable during or after a related failure (Figure 1). During modelling, probabilities are assessed to reflect experts' knowledge about strengths of the causal relationships. After an actually observed failure, the model can guide the user's process of collecting information about symptoms to distinguish faults.

The quality of fault diagnosis is limited by two factors: the availability of relevant information and our ability to draw conclusions that are justified by such information. Our choice of Bayesian networks as a formalism targets the second factor. However, employing this model can also contribute to the first factor, by focusing on what information is relevant, how to represent it, and how to obtain it to support automated fault diagnosis.

Our case study performed on the Condor middleware helped identify the following areas for future research:

- Selection of model variables

*This work is supported by the German Federal Ministry of Education and Research (BMBF) under grant No. 01C5968 and the German Research Foundation (DFG) under grant GRK 1076/1.

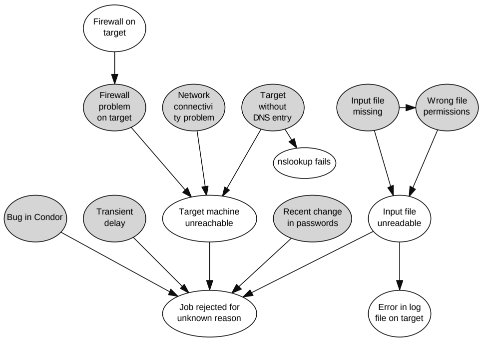

Figure 1: A Bayesian network for diagnosing rejected jobs in Condor.

- Representing object instances and states
- User interaction
- Costs of model construction and maintenance

Furthermore, we plan to develop a domain-specific vocabulary which can be used to describe common failure scenarios in Grid computing and to automate their diagnosis by incorporating available sources of information, such as distributed log files.

References

[BKM+04] Rob Barrett, Eser Kandogan, Paul P. Maglio, Eben M. Haber, Leila A. Takayama, and Madhu Prabaker. Field studies of computer system administrators: analysis of system management tools and practices. In *CSCW '04: Proceedings of the 2004 ACM conference on Computer supported cooperative work*, pages 388–395, New York, NY, USA, 2004. ACM Press.

[Pea88] Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers Inc., 1st edition, 1988.

[PH05] Jan Ploski and Wilhelm Hasselbring. The Callback Problem in Exception Handling. In Alexander Romanovsky, Christophe Dony, Jørgen L. Knudsen, and Anand Tripathi, editors, *Developing Systems that Handle Exceptions. Proceedings of ECOOP'05 Workshop on Exception Handling in Object-Oriented Systems*, pages 39–62. Department of Computer Science, LIRMM, University of Montpellier II, France, July 2005.

[RSB03] Joshua A. Redstone, Michael M. Swift, and Brian N. Bershad. Using Computers to Diagnose Computer Problems. In *Proceedings of the 9th Workshop on Hot Topics in Operating Systems*, 2003.

[TTL05] Douglas Thain, Todd Tannenbaum, and Miron Livny. Distributed computing in practice: the Condor experience: Research Articles. *Concurr. Comput.: Pract. Exper.*, 17(2-4):323–356, 2005.

[WIS06] WISENT. Wissensnetz Energiemeteorologie. `http://wisent.d-grid.de`, 2006. Retrieved: 2006-06-09.