

Using Adaptive Probing for Real-Time Problem Diagnosis in Distributed Computer Systems

I. Rish, M. Brodie, S. Ma, G. Grabarnik
 IBM T.J. Watson Research Center
 30 Saw Mill River Road, Hawthorne, NY 10532
 mbrodie, rish, shengma, genady@us.ibm.com

N. Odintsova
 Department of Computer Science
 Polytechnic University of New York
 nodintsova@hotmail.com

Abstract

In this work, we focus on cost-efficient techniques for real-time diagnosis in distributed systems that allow an adaptive, on-line selection and execution of appropriate measurements (tests). Particularly, one of our applications concerns fault diagnosis in distributed computer systems and networks by using test transactions, or *probes* (e.g., "traceroute" or "ping" commands). The key efficiency issues include both the cost of probing (e.g., the number of probes), and the computational complexity of diagnosis. In our past work (see (Rish, Brodie, & Ma 2002a)), we derived some theoretical conditions on the number of probes required for an asymptotic error-free diagnosis, and developed efficient search techniques for probe set selection that can greatly reduce the probe set size while maintaining its diagnostic capability (Brodie, Rish, & Ma 2001). Next, we considered the problem of real-time diagnosis as a probabilistic inference in Bayesian networks and investigated simple and efficient *local* approximation techniques, based on variable-elimination (the minibucket scheme (Dechter & Rish 2002)). Our empirical studies show that these approximations "degrade gracefully" with noise and often yield an optimal solution when noise is low enough, and our initial theoretical analysis explains this behavior for the simplest (greedy) approximation (Rish, Brodie, & Ma 2002a; 2002b). Our future work will focus on adapting more sophisticated approximation techniques, such as Generalized Belief Propagation (Yedidia, Freeman, & Weiss 2001), to real-time scenarios, and a real-time, incremental learning of Dynamic Bayesian Networks based on the historic data and the feedback on the diagnosis results.

Keywords: *Real-time Bayesian network inference and learning, active sampling, anytime approximate reasoning, resource-bounded computation.*

As distributed computer systems and networks continue to grow in size and complexity, tasks such as real-time fault localization and problem diagnosis become significantly more challenging. As a result, more sophisticated tools are needed that can assist in performing these management tasks by both responding quickly and accurately to the ever-increasing volume of system measurements, and also actively selecting minimum number of most-informative tests

to run. In other words, a "passive" data-mining approach must be replaced by active, real-time information-gathering and inference systems that can "ask right questions at the right time". Moreover, the focus on the cost-efficiency and scalability of real-time problem diagnosis is particularly important for making it useful in the context of extremely large geographically distributed (GRID) computing systems, and in the face of new technological challenges related to *autonomic computing*, an IBM's vision for a new-generation IT systems capable of self-management and self-repair.

We are currently working on creating a system for a real-time problem diagnosis by using *probing technology* which offers the opportunity to develop an approach to diagnosis that is more active than traditional "passive" event correlation (Kliger *et al.* 1997) and similar techniques. A probe is a command or a transaction (e.g., *ping* or *traceroute* command, an email message, or a web-page access request), sent from a particular machine called a *probing station* to a server or a network element in order to test a particular service (e.g., IP-connectivity, database- or web-access). A probe returns a set of measurements, such as response times, status code (OK/not OK), and so on. Probing technology is often used to measure the quality of network performance, often motivated by the requirements of service-level agreements (SLAs); however, using probing for real-time problem diagnosis appears to be an open area.




Figure 1: (a) A two-layer Bayesian network structure for a set $X = (X_1, X_2, X_3)$ of network elements and a set of probes $T = (T_1, T_2)$, and (b) its extension to a Dynamic Bayesian Network.

To use probes, probing stations must first be selected at

one or more locations in the network. Then the probes must be configured; it must be decided which network elements to target and which station each probe should originate from. Using probes imposes a cost, both because of the additional network load that their use entails and also because the probe results must be collected, stored and analyzed. Cost-effective diagnosis requires a small probe set, yet the probe set must also provide wide coverage, in order to locate problems anywhere in the network. By reasoning about the interactions among the probe paths, we construct an estimate of the *information gain* provided by each probe, and use this estimate as a probe selection heuristic. This yields a quadratic-time greedy-search algorithm which finds near-optimal probe sets. We also implement a linear-time algorithm which can be used to find small probe sets very quickly; a reduction of almost 50% in the probe set size is achieved. The results are reported in (Brodie, Rish, & Ma 2001; Rish, Brodie, & Ma 2002b). Moreover, in (Rish, Brodie, & Ma 2002a) we provide some theoretical bounds on the diagnosis error and derive necessary conditions on the number of probes required for an asymptotically error-free diagnosis.

Once the probes have been selected and sent, fault diagnosis is performed by analyzing the probe outcomes. In real-life scenarios this must be done in an environment of noise and uncertainty. For example, a probe can fail even though all the nodes it goes through are OK (e.g., due to packet loss). Conversely, there is a chance that a probe succeeds even if a node on its path has failed (e.g., dynamic routing may result in the probe following a different path). Thus the task is to determine the *most likely* configuration of the states of the network elements.

We use the graphical framework of Bayesian networks (Pearl 1988) that provides both a compact factorized representation for multivariate probabilistic distributions as well as a convenient tool for probabilistic inference. An example of a simple Bayesian network (BN) for problem diagnosis is shown in Figure 1a: a is bipartite (two-layer) graph where the top-layer nodes represent marginally independent faults or other problems (if the problems are not marginally independent, appropriate edges must be added between them) and the bottom-layer nodes represent probe results. In order to represent temporal dependencies, such network can be extended to a *k*-slice *Dynamic Bayesian Network* where each time-slice contains a copy of the above BN, and interslice dependencies are encoded by transition probabilities, as shown in Figure 1b (see (Rish, Grabarnik, & Odintsova 2002) for details).

Since the exact inference in large Bayesian networks is often intractable (NP-hard), we investigated the applicability of approximation techniques. The complexity of inference is usually associated with large probabilistic dependencies recorded during inference (clique size, or *induced width*) (Dechter 1996). Thus, a popular approximation approach is to restrict the complexity by focusing only on local interactions. We investigated the performance of two *local inference* techniques, *greedy-mpe* and *approx-mpe(1)*, which are the simplest members of the parametric family of variable-elimination algorithms known as *mini-bucket* ap-

proximations (Dechter & Rish 2002). Our empirical studies show that these approximations "degrade gracefully" with noise and often yield an optimal solution when noise is low enough, and our initial theoretical analysis explains this behavior for the simplest (greedy) approximation (Rish, Brodie, & Ma 2002a; 2002b). The mini-bucket scheme is closely related to other local approximations, such as *iterative belief propagation (IBP)* and *generalized belief propagation (GBP)* algorithms (Yedidia, Freeman, & Weiss 2001). We plan to investigate those approximation approaches in our future work and hope to extend our analysis of the mini-bucket scheme that can be viewed as simplified, non-iterative version of those state-of-the-art approaches.

Finally, there are several directions for future work we are planning to pursue:

- adaptive probing - i.e. adjusting the probe set dynamically in response to the state of the network;
- extending local approximation techniques (mini-buckets, GBP) to incremental, real-time scenarios;
- handling intermittent failures, dynamic routing, and other non-stationarities in the network state and behavior; the latter would require tuning the model, i.e. online learning, which can be also done *actively* by using probe selection to improve learning (e.g., only update the part of the model which is currently relevant).

References

- Brodie, M.; Rish, I.; and Ma, S. 2001. Optimizing probe selection for fault localization. In *Distributed Systems Operation and Management*.
- Dechter, R., and Rish, I. 2002. Mini-buckets: A General Scheme for Approximating Inference. *To appear in J. of ACM*.
- Dechter, R. 1996. Bucket elimination: A unifying framework for probabilistic inference. In *Proc. Twelfth Conf. on Uncertainty in Artificial Intelligence*, 211–219.
- Kliger, S.; Yemini, S.; Yemini, Y.; Ohsie, D.; and Stolfo, S. 1997. A coding approach to event correlation. In *Intelligent Network Management (IM)*.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.
- Rish, I.; Brodie, M.; and Ma, S. 2002a. Accuracy vs. Efficiency Trade-offs in Probabilistic Diagnosis. In *To appear in the Proceedings of the The Eighteenth National Conference on Artificial Intelligence (AAAI2002)*, Edmonton, Alberta, Canada.
- Rish, I.; Brodie, M.; and Ma, S. 2002b. Intelligent probing: a Cost-Efficient Approach to Fault Diagnosis in Computer Networks. *Submitted to IBM Systems Journal*.
- Rish, I.; Grabarnik, G.; and Odintsova, N. 2002. Dynamic Bayesian Network Approach to Real-time Problem Diagnosis using Probing. Technical report, IBM T.J. Watson Research Center.
- Yedidia, J.; Freeman, W. T.; and Weiss, Y. 2001. Generalized belief propagation. In *NIPS 13*, 689–695. MIT Press.