1995 Pacific Rim International Symposium on Fault-Tolerant Systems, Newport Beach, CA, Dec. 1995, pp. 190-191.

The Need for Measurement Based Dependability Evaluation

Myron Hecht and Dong Tang
SoHaR Incorporated, Beverly Hills, California

Abstract

Defensible quantitative assessments of the dependability of fault tolerant systems should be a research priority. This position paper characterizes the need for quantitative empirically based dependability assessment, describes some of the previous work in this area, and identifies problems.

Introduction

The major benefit of fault tolerance is high dependability. However, achieving these benefits requires a defensible quantitative characterization of dependability. For example, International Joint Aviation Regulations impose maximum acceptable probabilities for failures of systems in passenger transport aircraft. The U.S. Code of Federal Regulations also establishes maximum acceptable probabilities for radioactive releases from nuclear power plants. The challenge faced by the developers of fault tolerant systems for these applications need to show that such quantitative requirements are being met. More detailed quantitative characterizations can identify system bottlenecks, and providing insight for decision making.

Estimation of System-Level Reliability and Availability

Quantitative characterization of system level reliability and availability can be achieved by creating models and using measurement-based parameters in these models [Tang95]. Statistical estimation of reliability and availability parameters and reliability modeling based on these parameters has been a research field in computer engineering for 15 years [Iyer93]. These analyses are based on operational logs and failure data.

In the process of collecting and analyzing such data, additional studies can done undertaken for more detailed examinations of underlying causes. For example, analyses of workload and failure data collected from IBM mainframes [Butner80] and DEC minicomputers [Castillo81] revealed that the average system failure rate is strongly correlated with the average workload on the system. Recent studies of data from DEC [Tang92] and Tandem [Lee91] systems showed that correlated failures across processors are significant in multicomputers and their impact on dependability is significant.

Estimation of Software Reliability

A defensible characterization of software failure rates essential to a system-level characterization of reliability and availability. Unfortunately, no generally accepted methodology of software reliability prediction has emerged.

While software reliability PREDICTION is unproven, reliability MEASUREMENT is not. Reliability life testing and measurement have a long and established history in the nuclear, aerospace, and DoD communities. In these measurement-based approaches, digital systems are regarded as consisting of components (equipments) whose operating times, failure rates, correlated failure probabilities, and recovery probabilities can be measured. Similar approaches have also been used to measure reliability of operational software for several commercial operating systems [Lee93].

The underlying assumption in these measurement-based approaches is that the fundamental failure mechanisms are triggered stochastically, i.e., are non-deterministic ("Heisenbugs"). However, an independent class of failures is related to the software successfully running to completion, but producing an unacceptable output. For example, an electronic speed control on a turbine may in fact not shut down the device in an over speed condition even though there was no crash, hang, stop, or delay failure. This deterministic failure condition may be traced to a logic fault in the code or an incorrect set of parameters (e.g., the RPM threshold for that particular turbine under the

specified set of pressures and temperatures). However, the root cause of the failure may in fact lie much deeper, i.e., defects in the system requirements or software requirements.

The techniques and methodologies for estimating the probabilities for these deterministic incorrect response failures are very immature. It is tempting to "wish them away" by positing that an adequate V&V or integration testing program should uncover them. However, resources are finite, and it is rarely feasible to provide sufficient time or money to perform the level of testing needed to uncover all such failures, even in systems designed for high dependability.

From a practical perspective, when estimating software failure rates, one should look not only at failures that cause losses or delays of system services (e.g., crash, hang, stop) but also incorrect response failures. If the proportion of incorrect responses at the final stages of testing or integration, or in initial operation, then reliability predictions made exclusively on the basis of stochastic failures may not be valid.

Data Collection

Obtaining adequate data from which to assess reliability and availability is critical to any measurement-based methodology. This obvious principle can be difficult to implement in practice for dependability assessments because of the constraints of an expensive testing program or impending project deadlines. Adequate data means monitoring and recording events of interest such as failures and recoveries of components, as well as performance parameters of the target system while it is operating under representative workloads. It also means collecting data on failure modes so that an assessment of the importance of deterministic failures can be made. The events and parameters to be collected should be representative of the system operation and meaningful for the assessment of the system. Measurements should be made continuously for a sufficient period to yield statistically significant data. Operating logs should include information about the location, time, and type of the error, the system state at the time of failure or abnormal operation, and error recovery (e.g., retry) information where applicable.

Conclusion

While there is still ongoing research in measurement-based analysis of computer system dependability, the techniques developed in the area have achieved significant experimental results. Measurement-based analysis can also provide verification of assumptions and parameters used in the design models. The results are useful for designing and maintaining highly dependable computer systems intended for use in critical applications such as flight control, ground transportation, air traffic control, and nuclear power plant safety systems.

References

[Butner80] S.E. Butner and R.K. Iyer, "A Statistical Study of Reliability and System Load at SLAC," Proc. 10th Int. Symp. Fault-Tolerant Computing, pp. 207-209, Oct. 1980.

[Castillo81], X. Castillo and D.P. Siewiorek, "Workload, Performance, and Reliability of Digital Computer Systems," *Proc. 11th Int. Symp. Fault-Tolerant Computing*, pp. 84-89, July 1981.

[Iyer93] R.K. Iyer and D. Tang, "Experimental Analysis of Computer System Dependability", Technical Report CRHC-93-15, Center for Reliable and High-Performance Computing, University of Illinois at Urbana-Champaign, July 1993.

[Lee93] I. Lee, D. Tang, R.K. Iyer, and M.C. Hsueh, "Measurement-Based Evaluation of Operating System Fault Tolerance," *IIEEE Transactions on Reliability*, pp. 238-249, June 1993.

[Tang92] D. Tang and R.K. Iyer, "Analysis and Modeling of Correlated Failures in Multicomputer Systems," *IIEEE Trans. Computers* Vol. 41, No. 5, pp. 567-577, May 1992.

[Tang95] D. Tang and M. Hecht, "Evaluation of Software Dependability Based on Stability Test Data", *Proc. 11th Int. Symp Fault-Tolerant Computing*, Pasadena, California, June, 1995