EXAMINING COINCIDENT FAILURES AND USAGE-PROFILES IN RELIABILITY ANALYSIS OF AN EMBEDDED VEHICLE SUB-SYSTEM

Frederick T. Sheldon and Kshamta Jerath
Software Engineering for Dependable Systems Laboratory©
School of EECS, Washington State University
Pullman, Washington 99164-2752, USA
E-mail: sheldon@acm.org | kjerath@eecs.wsu.edu

Stefan A. Greiner
Research Information and Communication
DaimlerChrysler, HPC T728 Stuttgart 70546, Germany.
E-mail: stefan.greiner@daimlerchrysler.com

KEYWORDS

Reliability, Stochastic modeling, Coincident-failures, Correlated failures, Usage-Profiles, Operational profiles.

ABSTRACT

Structured models of systems allow us to determine their reliability, yet there are numerous challenges that need to be overcome to obtain meaningful results. This paper reports the results and approach used to model and analyze the Anti-lock Braking System of a passenger vehicle using Stochastic Petri Nets. Special emphasis is laid on modeling extra-functional characteristics like coincident failures among components, severity of failure and usage-profiles of the system. Components generally interact with each other during operation, and a faulty component can affect the probability of failure of other components. The severity of a failure also has an impact on the operation of the system, as does the usage profile - failures which occur during active use of the system are the only failures considered (i.e., in reliability calculations).

INTRODUCTION

A complex system (like an embedded vehicle system) is composed of numerous components and the probability that the system survives (efficient or acceptable degraded operation) depends directly on each of the constituent components. The reliability analysis of a vehicle system can provide an understanding about the likelihood of failures occurring in the system and an increased insight to manufacturers about inherent "weaknesses." (Jerath and Sheldon 2001) Therefore, this paper provides the basis of how to gain such insights using SPNs, for an example, analyzing an ABS system. Special emphasis is laid on modeling extra-functional characteristics like coincident failures among components, severity of failure and usage-profiles of the system.

If a system does not contain any redundancy – that is, if every component must function properly for the system to work – and if component failures are statistically independent, then the system reliability is simply the product of the component reliabilities. Furthermore, the failure rate of the system is the sum of the failure rates of the individual components (Siewiorek and Swarz 1992). The assumption that failures occur independently (in a statistical sense) in hardware components is a widely used model for predicting the reliability of hardware devices. However, components generally interact with each other during operation, and a faulty component can affect the probability of failure of other components too (Balbo 2000). Such failures are not **coincident** in the sense that they occur simultaneously, but in the fact that failure of one increases the probability of the failure of another.

Another aspect of modeling failures occurring in the system is their **severity**. Severity of a failure is the impact it has on the operation of the system. It is closely related to the threat the problem poses, in functional terms, to the correct operation of the system (Vouk 2000). Severity is an important candidate to weight the data used in reliability calculations and should be incorporated into the model to determine the probability that the system survives, including efficient or acceptable degraded operation.

The reliability of a system also depends on its **usage profile** – users interact with the system in an intermittent fashion, resulting in operational workload profiles that alternate between periods of "Active" and "Passive" use. Reliability is concerned with the service that is actually delivered by the system as opposed to a system's *capacity* to deliver such service (Meyer 2000). Specifically, while considering usage profiles, faults need not necessarily cause failures since they can be repaired; failures occurring during "active" use of the system *only* should contribute to reliability calculations.

In (Sheldon et al. 2000), the authors presented Stochastic Petri Net (SPN) models of a vehicle dynamic driving regulation (DDR) system. Subsystem representations of the Anti-lock Braking system (ABS), the Electronic Steering Assistance (ESA), the traction control (TC) and a combined model were developed and analyzed for critical failures. In this paper, we focus on the Anti-lock Braking system and develop Stochastic Petri Net models to model the coincident failures of components, severity of failures and usage-profiles. Naturally this is but one component of the total system and the issue of scalability of this approach is a subject for future work.

The paper is organized as follows. The section "System Description and Modeling Approach" briefly describes the structural and functional aspects of an Anti-lock Braking System (ABS) and the Petri Net approach to modeling. Sections "Modeling Coincident Failures and Severity" and "Modeling Usage-Profiles" present the assumptions, SPN models and results for the Petri-nets modeling coincident failures and severity of failures, and usage-profiles respectively. Finally, the conclusion and the scope for future work are discussed in the section "Conclusion and Future Work".

SYSTEM DESCRIPTION AND MODELING APPROACH

In this section, we briefly examine the structural composition of an Anti-lock Braking System and its functionality. Stochastic Petri Nets (SPNs) were used to model the system and the Stochastic Petri Net Package (SPNP) to analyze the models. The modeling and analysis approach is discussed later in this section.

Anti-lock Braking System

The Anti-lock Braking System is an integral part of the total braking system in a vehicle. Applying excessive pressure on the brake pedal, or panic slamming the brake pedal, can cause wheels to lock up and possibly send the vehicle careening into a terrifying skid. The ABS prevents wheel lockup during an emergency stop by modulating the brake pressure and permits the driver to maintain steering control while braking.

The ABS consists of the following major components:

- Wheel Speed Sensors: These measure wheel-speed and transmit information to an electronic control unit.
- Electronic Control Unit (Controller): This receives information from the sensors, determines when a wheel is about to lock up and controls the hydraulic control unit.
- Hydraulic Control Unit (Hydraulic Pump): This controls the pressure in the brake lines of the vehicle.
- Valves: Valves are present in the brake line of each brake and are controlled by the hydraulic control unit to regulate the pressure in the brake lines.

Under braking, the electronic control unit (ECU) “reads” signals from electronic sensors monitoring wheel rotation. If a wheel’s rate of rotation suddenly decreases, the ECU orders the hydraulic control unit (HCU) to reduce the line pressure to that wheel’s brake. Once the wheel resumes normal operation, the controls restore pressure to its brake. Depending on the system, this cycle of “pumping” can occur at up to 15 times per second. The result is that the tire slows down at the same rate as the car, with the brakes keeping the tires very near the point at which they will start to lock up. This gives the system the highest steering capability. Anti-lock braking systems use different schemes depending on the type of brake in use. In the model developed we assume a four channel four sensor ABS. The model can be easily modified to represent other ABS schemes.

Modeling And Analysis Using SPNs

A powerful tool for modeling systems composed of several processes (such as a failure process and a repair process) is the Markov Model. Markov Models are a basic tool for both reliability and availability modeling. Stochastic Petri Nets (SPN) can be used to generate the (large) underlying Markov chain automatically starting from a concise description of the system. Stochastic Petri Nets are commonly used to evaluate the performance and reliability of complex systems (Balbo 2001) because the graphical nature of SPNs lends itself to a more intuitive understanding of the system’s inner workings and allows one to understand dependencies better. This enables one to identify conflicts and address localities where the overall system performance is more significantly affected.

Since the system we study here is very complex, this prevents us from making a direct analysis. Two distinct problems that arise while using SPNs are largeness and stiffness (Popstojanova and Trivedi 2000). The size of a Markov Model for the evaluation of a system grows exponentially with the number of components in the system and stiffness is due to the different orders of magnitude between the rates of failure-related events in different components. A series of abstraction steps are needed and the key element in our modeling approach was to identify the essential components of the system, the different ways in which they interact and introduce various assumptions. The details of the models developed and the assumptions made are discussed in the next two sections.

We described the models in CSPL (C-based Stochastic Petri net Language) and the stochastic analysis was carried out using SPNP (Stochastic Petri Net Package), a versatile modeling tool (Ciardo et al. 1989). The models were solved using Version 6 of SPNP installed on a Sun Ultra 10 (400Mhz) with 500MB of memory (dedicated to solving the models). The models took approximately 5 days of continuous execution before converging to solution. This time may have been drastically reduced we believe had the Multi-level solution method been available within the SPNP package (Greiner and Horton 1996).

MODELING COINCIDENT FAILURES AND SEVERITY

The assumption that failures occur independently is a widely used and often successful model for predicting the reliability of hardware devices. However, as mentioned above, components generally interact with each other during operation, and a faulty component can affect the probability of failure of other components too (Balbo 2000). Severity of a failure is the impact it has on the operation of the system and is an important candidate to weight the data used in reliability calculations. In this section, we describe the Petri net models developed to model coincident failures and severity of failures for the Anti-lock Braking System.

Assumptions

To allow a Markov chain analysis, the time to failure of all components is assumed to have an exponential distribution. This signifies that the distribution of the remaining life of a component does not depend on how long the component has been operating. The component does not “age” or it forgets how long it has been operating, and its eventual breakdown is the result of some suddenly appearing failure, not of gradual deterioration (Trivedi 2001). While this might be true for electronic components, the failure of other mechanical parts like valves might occur due to gradual deterioration. However, mechanical parts are generally replaced at regular intervals and essentially can be assumed not to age for our purposes. Hence, the assumption of an exponential distribution of failures for all components is justified.

To consider the severity of failures, every component is assumed to operate in three modes: normal operation, degraded operation or causing loss of stability. The system is assumed to ***fail*** when more than five components function in a degraded state or, more than three components cause loss of stability; or the failure of an important component causes the loss of the vehicle. A component operating in a degraded condition causes its failure rate to increase by two orders of magnitude, while a component causing loss of stability causes the failure rate to increase by four orders of magnitude. The correlation between failure rates of two "related" components (to model coincident failures) is consistent with the above scheme.

Since the model is an abstraction of a real world problem, predictions based on the model must be validated against actual measurements collected from the real phenomena. A poor validation may suggest modifications to the original model (Trivedi 2001).

Model

The ABS is represented as a combination of all the important components it consists of, as shown in Figure 1. It represents the operation of the ABS under normal, degraded and lost stability conditions. Loss of vehicle, extreme degraded operation and extreme loss of stability signify critical failures and determine the halting condition for the model. The model is initialized with a single token in the start place. When the central_op and the axle_op transitions fire, a token is deposited in each place that represents a component of the ABS. The operation of each component is now independent of every other component (except where coincident failures are modeled explicitly).

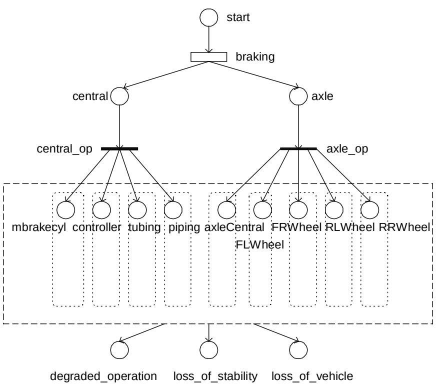

Figure 1: The ABS Model

The model of a component of the ABS is shown in Figure 2. The component depicted here is the controller. Every component either functions "normally" as shown by the controllerOp transition or "fails" as shown by the controllerFail transition. A failed component may either cause degraded operation, loss of stability or loss of vehicle. The probability of any one of these three transitions occurring is different for each component. When the failure causes either degraded operation or loss of stability, the component continues to operate, though the failure rate increases by two and four orders of magnitude respectively. At the same time, a token is deposited in either controllerDegraded or controllerLOS, as the case may be, to be able to identify which component caused the entire system to fail.

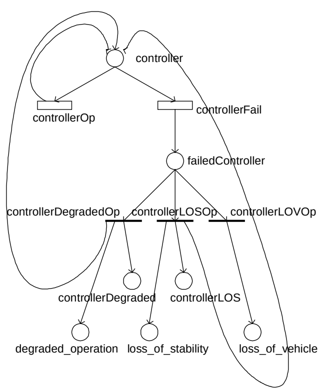

Figure 2: The Model of a Component

Coincident failures are modeled in a similar manner. The rule for calculating failure rates is shown in Figure 3. The failure of a component A to a degraded mode causes the failure rate of a "related" component B to increase by two orders of magnitude. The failure of component A to a lost-stability mode causes the failure rate of a "related" component B to increase by four orders of magnitude.

```
function failureRateForB()
{
    // other calculations for severity of failure

    // coincident failures
      if failedA(degraded) then
                failureB = failureB * 100;
      else if failedA(loss of stability) then
                failureB = failureB * 10000;
}
```

Figure 3: Rule for Failure Rates

The function that calculates the failure rate of the transition *controllerFail* is shown in Figure 4. It is assumed that tubing malfunction affects the operation of the controller. Hence, while calculating the failure rate of the controller, the normal rate is increased by two orders of magnitude if the tubing has failed causing degraded operation (indicated by a token in the *tubingDegraded* place).

```
double controllerRate()
{
    double controller_rate = 0.0000006;

    if (mark("controllerLOS") > 0) return controller_rate * 10000;
    if ((mark("controllerDegraded") > 0)
                    || (mark("tubingDegraded") > 0))
         return controller_rate * 100;
    return controller_rate;
}
```

Figure 4: Variable Rate to Model Coincident Failures

Only a few coincident failures have been represented in the model. However, coincident failures between other components can be easily modeled by suitably modifying the failure rate function of the component in question using the rule shown in Figure 3. The model is easily extensible to include other components deemed relevant to the ABS.

Results And Discussion

The Stochastic Petri Net Package (SPNP) allows the computation of steady state, transient, cumulative, time-averaged, "up-to-absorption" measures and sensitivities of these measures (Muppala et al. 1994). Transient analysis of the ABS model was carried out and the reliability was measured between 0 and 50K hours. The transient analysis duration of the models developed was deliberately conservative - the average life span of a passenger vehicle ranges from 3000 – 9000 hours (Essentially the average hours of operation for a passenger vehicle per year range from 300-600 hours/year and the average lifetime is 10-15 years). The expected values of reliability at various time instances were determined and plotted as a function of time.

In Figure 5, the Y-axis gives the measure of interest - the reliability; while the time range (0 to 50K hours) is shown along the X-axis. The shape of the curve is not a property of the system but of how the data was collected from the Petri net model. As expected, the reliability steadily decreases with time. The line marked with diamonds indicates the reliability function when coincident failures are modeled and the line marked with squares indicates the reliability function when coincident failures are not modeled (entirely overlapped). For the limited number of coincident failures that were modeled, it is clear that the Mean Time to Failure (MTTF) for the model with coincident failures (784,856.4 hrs) is approximately 421 hours less than the model without coincident failures (785,277.6 hrs).

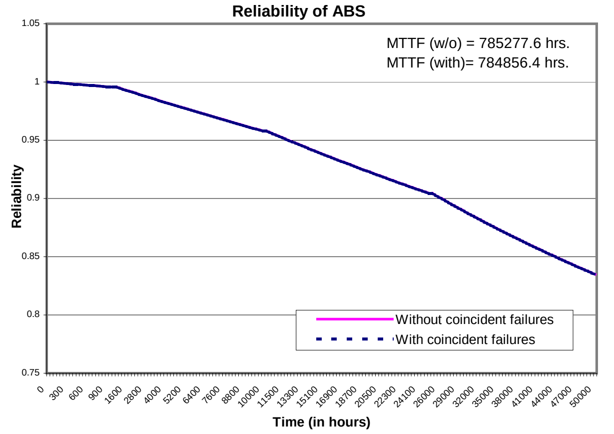

Figure 5: Reliability Analysis Results for Coincident Failures

Figure 6 displays the difference between the two reliability functions more subtly. The reliability functions diverge starting around 350 hours of operation, and the difference becomes discernible after around 13K hours of operation. The difference continues to increase with time. It is significant to note that the difference in Mean Time To Failure between the two cases becomes marked only beyond the average lifetime of the vehicle. For the limited number of coincident failures that have been modeled, the difference of 421 hours in the two cases is considered well within the confidence interval. However, it is

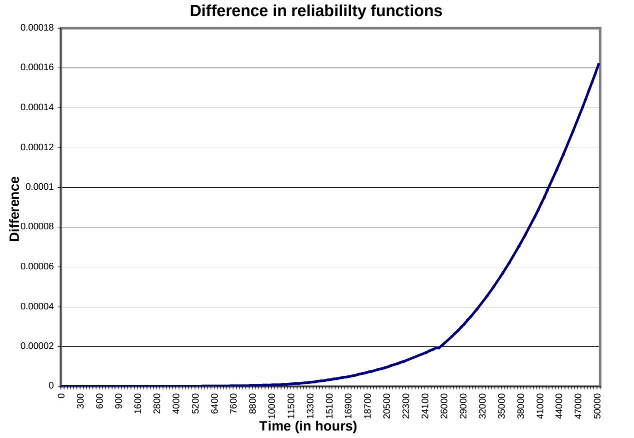

Figure 6: Difference in Reliability Functions

evident that the model representing the coincident failures predicts the system reliability closer to the real picture.

MODELING USAGE-PROFILES

A software-based product's reliability depends on just how a customer will use it. The operational profile – quantitative characterization of how a system will be used – is essential in software reliability engineering (Musa 1993). We extend the idea of operational profiles – considering the use of a software system during testing; into usage profiles – the usage of the system (hardware and software) for modeling and reliability analysis. The usage profile considers the intermittent use of a system – alternate periods of active and passive use. Such intermittent use influences the mean time to failure and reliability of the system (Meyer 2000). In this section, we describe the Petri net models developed to model usage-profiles for the Anti-lock Braking System.

Assumptions

Unlike traditional reliability models where repair of components is not considered, when considering intermittent use it is important to note that faults need not necessarily cause failures. Faults occurring only during the active use cause failures while those occurring during passive use can be repaired. Hence repair can affect reliability calculations. For simplicity, we assume an infinite repair rate of all components.

Further, in order to comprehend the significance of intermittent use on reliability, we assume two usage-profiles exceedingly different in degree. The first profile models sparse use of the Anti-lock Braking System e.g. a driver who is extremely cautious while driving the vehicle (longer periods of passive use). The second usage profile models dense use of the anti-lock braking system e.g. a driver in perilous conditions like driving over ice (frequent active use periods).

Again, for simplicity and to allow Markovian analysis, the active period duration is assumed to be exponentially distributed, as are the failure rates of the components. The second usage-profile is assumed to have a rate two orders of magnitude greater than the first usage profile. In order to work around the stiffness problem in Petri nets caused by the difference in magnitude between the failure rates of the components and the active period duration distribution rates, the duration distribution rates are assumed to be factored by the failure rates of individual components.

Model

To incorporate the usage-profiles scenario in the ABS model, the model of each individual component as depicted in Figure 2 was extended as shown in Figure 7. The figure again shows the controller component with the additions to the model marked in red. When a failure occurs (*failedController*) the next step is to determine whether the system was in active use or not, because only a failure occurring when the system is in active use contributes towards reliability calculations. The parameter 1/mu indicates the mean duration of active use while the parameter 1/alpha indicates the mean duration of passive use.

If the failure occurs during the active-use period (*inUseController*), the system either continues to operate in a degraded mode (degraded operation - *controllerDegradedOp*

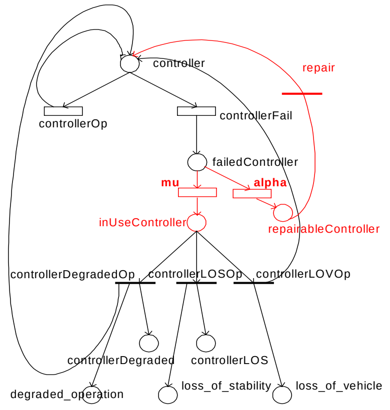

Figure 7: SPN Model with Usage Parameters

or lost-stability - *controllerLOSOp*) or causes loss of vehicle (*controllerLOVOp*). If the failure occurs during passive use of the system (*repairableController*), the fault can be repaired and an infinite repair rate is assumed. The system continues to operate as if no failure had occurred.

To work around the state explosion problem that occurred due to the increase in the number of states in the model, the model was simplified. The usage parameters were incorporated while calculating the failure rate itself for each component. The modified function for calculating the failure rate in light of the usage-profile is shown in Figure 8. The value of *mu* was assumed to be 2.5 for infrequent active use periods and 250 for frequent active use periods. As stated in the assumptions and shown in Figure 8, the value of these usage distributions was factored by the actual failure rate of the component to avoid stiffness in the model.

```
double controllerRate()
{
        double controller_rate = 0.0000006;

        // usage parameter
        controller_rate += controller_rate * mu();

        if (mark("controllerLOS") > 0) return controller_rate * 10000;
        if ((mark("controllerDegraded") > 0)
                        || (mark("tubingDegraded") > 0))
                return controller_rate * 100;
        return controller_rate;
}
```

Figure 8: Variable Rate to Model Usage Parameter

Results And Discussion

Transient analysis of the developed ABS model was carried out and the reliability was measured between 0 and 50K hours. The expected values of reliability at various time instances and different usage profiles was determined and plotted as a function of time. The results are depicted in Figure 9. The Y-axis gives the measure of interest - the reliability; while the time range (0 to 50K hours) is shown along the X-axis. The shape of the curve is not a property of the system but of how the data was collected from the Petri net model.

As expected, the reliability steadily decreases with time. The top line indicates the reliability function when the usage of the system is infrequent and the bottom line indicates the reliability function when the usage of the system is frequent. The reliability of the system with heavy usage decreases alarmingly within the first 1K hours of operation, while the reliability of

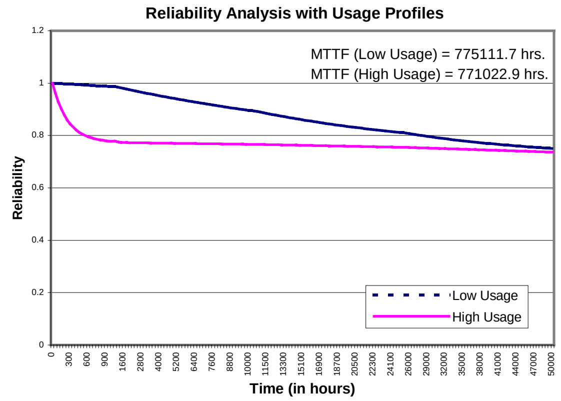

Figure 9: Reliability Analysis Results for Usage Profiles

the system with not so heavy usage decreases perceptibly only after 2.5K hours of operation and then steadily afterwards. Also, the mean time to failure (MTTF) for the high usage case is 771022.9 hours as opposed to 775111.7 hours for the low usage case, a difference of approximately 4089 hours.

An important fact to consider is that some components are used only for a few minutes during the entire lifetime of the vehicle (10-15 years) while other components like the tubing are used all of the time during that period. Hence, the usage of different components is different even within a given usage profile and will affect the actual reliability. However, what is important is the approach we used and the results clearly indicate that it is important to consider the usage profiles while determining the reliability for any given system.

CONCLUSION AND FUTURE WORK

In this paper, we have shown how to model coincident failures, severity and usage-profiles in the Anti-lock Braking system of a passenger vehicle using Stochastic Reward Nets. We made some simplifying assumptions in order to manage the complexity of the system being modeled apart from handling the general challenges in the modeling like state explosion and stiffness. The Stochastic Petri Net models were developed for a four channel four sensor ABS. The model, however, is easily extensible to model other schemes of ABS. Other coincident failures between components can be easily modeled by suitably modifying the failure rate function of the component in question. Similarly, other profiles with different usage parameters can be easily incorporated and analyzed. SPNP was used to specify the system and carry out the reliability analysis.

A major obstacle in modeling using SPNs was the persistent state explosion problem. This caused the programs to abort due to insufficient memory while solving the Markov chains. Stochastic Activity Networks (SANs) (Sanders and Meyer 2001) are a stochastic extension to SPNs and are used for performability evaluation. As shown in Figure 10, composed models in SANs exploit symmetries in the model to reduce the number of reachable states. The models can be specified and analyzed using UltraSAN, a software tool for model-based performance, dependability and performability evaluation of computer, communication and other systems (Sanders 1994-95). The goal of future work is to specify SAN models for the Anti-lock braking system, analyze them using UltraSAN and compare the results obtained for SPN models.

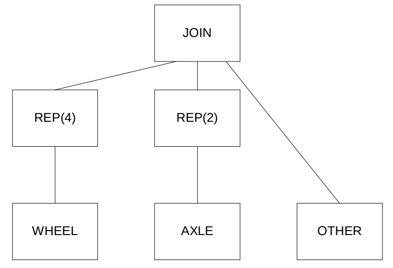

Figure 10: Composed SAN Model

Further, the Anti-lock Braking system is a small part of the DDR (Dynamic Driving Regulation) system which consists of subsystems like the Anti-lock Braking system (ABS), the Electronic Steering Assistance (ESA), the traction control (TC) (Sheldon et al. 2000). Another goal is to develop and analyze a model extended for critical (mission-safety) failures that scales well for the composed model emphasizing coincident failures, severity of failures and usage-profiles and analyze it for critical failures.

REFERENCES

Balbo, G., 2000. Professor, Universita di Torino, Italy. Personal Communication at EEF-Summerschool Formal Methods and Performance Analysis, Netherlands, July 2000.

Balbo, G., 2001. "Introduction to Stochastic Petri Nets", *Lecture Notes in Computer Science*, No.2090, 84-155.

Ciardo, G., Muppala, J.K. and Trivedi, K., 1989. "SPNP: Stochastic Petri Net Package", In *Proceedings of the Intl. Workshop on Petri Nets and Performance Models*, Kyoto, Japan, IEEE Computer Society Press. 142-151.

Greiner, S. and Horton, G., 1996. "Analysis of Stiff Markov Chains with the Multi-level Method", In *Proceedings of the European Simulation Symposium, ESS '96*, 801-805.

Jerath, K. and Sheldon, F.T., 2001. "Reliability Analysis of an Anti-lock Braking System using Stochastic Petri Nets", In *Proceedings of the PMCCS5*, Erlangen, Germany, Springer Verlag. 56-60.

Meyer, J., 2000. Professor, University of Michigan, Ann Arbor, MI. Personal Communication at PMCCS5, Erlangen, Germany, September 2001.

Muppala, J.K., Ciardo, G. and Trivedi, K., 1994. "Stochastic Reward Nets for Reliability Prediction", *Communications in Reliability, Maintainability and Serviceability*, No.1(2), 9-20.

Musa, J.D., 1993. "Operational Profiles in Software-Reliability Engineering", *IEEE Software*, No.10(2), 14-32.

Popstojanova, K.G. and Trivedi, K., 2000. "Stochastic Modeling Formalisms for Dependability, Performance and Performability", *Lecture Notes in Computer Science*, No.1769, 403-422.

Sanders, W.H., 1994-95. *UltraSAN User's Manual version 3.0*, http://www.crhc.uiuc.edu/PERFORM/Papers/USAN_papers/manual_v3.0_all.pdf.

Sanders, W.H. and Meyer, J., 2001. "Stochastic Activity Networks: Formal Definitions and Concepts", *Lecture Notes in Computer Science*, No.2090, 315-343.

Sheldon, F.T., Greiner, S. and Benzinger, M., 2000. "Specification, Safety and Reliability Analysis Using Stochastic Petri Net Models", In *Proceedings of the Tenth International Workshop on Software Specification and Design*, San Diego, California, IEEE Computer Society. 123-132.

Siewiorek, D.P. and Swarz, R.S., 1992. *Reliable Computer Systems: Design and Evaluation*, Digital Press, Burlington, MA.

Trivedi, K., 2001. *Probability and Statistics with Reliability, Queuing and Computer Science Applications*, John Wiley & Sons, New York.

Vouk, M.A., 2000. "Software Reliability Engineering", In *Proceedings of the 2000 Annual RELIABILITY and MAINTAINABILITY Symposium*, Los Angeles, CA, IEEE Computer Society.

AUTHOR BIOGRAPHIES

Frederick T. Sheldon is an Assistant Professor at the Washington State University teaching and conducting research in the area of software engineering. His research is concerned with developing and validating methods and supporting tools for the creation of safe and correct software. Recent studies conducted at the SEDS laboratory (Software Engineering for Dependable Systems) have focused on verification and validation of systems using modeling and analysis of both logical and stochastic properties. Dr. Sheldon received his Ph.D. at the University of Texas at Arlington (UTA) and has worked at NASA Langley and Ames Research Centers in various capacities since 1993. Prior to that, he worked as a Software Engineer in the area of avionics and diagnostics software development for the YF-22, F-16 and Tornado aircraft programs at General Dynamics and Texas Instruments. He is a member of the IEEE Computer and Reliability Societies, ACM, AIAA, and the Tau Beta Pi and Upsilon Pi Epsilon. www.eecs.wsu.edu/~sheldon | sheldon@acm.org

Kshamta Jerath is a Graduate Student at the Washington State University pursuing a master's degree in Computer Science. Her research interests include Software Engineering and Formal Methods (Stochastic Petri Nets) used in safety and reliability analysis. She is currently working as a Teaching Assistant for the Introduction to Computer Networks course at WSU. Ms. Jerath holds a bachelor's degree in Computer Engineering from Delhi College of Engineering, India and is a Sun Certified Java 2 Programmer. She was working as a software engineer with IBM India till December 2000 in the area of web application development and three-tier client server applications. She is a member of the IEEE. www.eecs.wsu.edu/~kjerath | kjerath@eecs.wsu.edu

Stefan A. Greiner studied computer Science at the University of Erlangen-Nuremberg from 1991 - 1994. After completing in 1994 he spent two years at Duke University, Durham USA to extend his studies and write a book with the title "Queueing Networks and Markov Chains" together with Prof.Trivedi, Prof. DeMeer and Dr. Bolch. After returning to Erlangen in 1995 Dr. Greiner worked on his PhD with the title "Modeling and Analysis of Operating Systems Using Extended QN Techniques and Petri Nets". Since 1997, Dr. Greiner works at DaimlerChrysler research in the RIC/AS department and focuses in the areas of Quality and reliability analysis.

www4.informatik.uni-erlangen.de/~sngreine/ | stefan.greiner@daimlerchrysler.com