CRISP Seminar: Mar 3, 2021 at 12-1pm

Virtual room: https://engineering.purdue.edu/crisp/seminar

Nirmal Saxena
Distinguished Engineer, NVIDIA

Road to Resilient Autonomous Cars is Paved with Testability & Diverse Redundancy

According to the 2015 vehicular accidents report [www.nhtsa.gov], there were more than 35000 fatal crashes and more than 6 million non-fatal crashes. Translating the 35K fatal crashes, over 3 trillion driven miles, to a FIT (failures in time, time = one billion hours) rate we get a fatality FIT rate in the range 250-500. A drive system, in a fully autonomous car, that replaces the human driver must at least be an order of magnitude more resilient. In fact the ISO 26262 Auto Safety Standard, stipulates a probabilistic metric for hardware failures (PMHF) to be at most 10 FITs. This target is almost two orders of magnitude improvement over driver related fatal accident FIT rate; notwithstanding, the fact that not all hardware failures result in fatalities. Autonomous or self-driving car initiative is creating a new center stage for resilient computing and design for testability. This is very apparent from reading the ISO26262 specification, which is about the functional safety for automotive equipment applicable throughout the lifecycle of all automotive electronic and electrical safety-related systems. By way of a detailed review of the ISO26262 standard, this tutorial pays tribute to all of the important and significant ideas that have come through the past 40+ years of research in fault-tolerant computing and design-for-testability. One of the key use cases for self-driving car is neural-network based deep learning. Reliability models that explore the dynamic trade-offs between resiliency and performance requirements for deep learning are examined. Deep neural networks use the computational power of massively parallel processors in applications such as autonomous driving. Autonomous driving demands resiliency (as in safety and reliability) and trillions of operations per second of computing performance to process sensor data with

extreme accuracy. This talk examines various approaches to achieve resiliency in autonomous cars and makes the case for design diversity based redundancy as a viable solution.

Speaker Biography

Nirmal R. Saxena is currently a distinguished engineer at NVIDIA and is responsible for high-performance and automotive resilient computing. From 2011 through 2015, Nirmal was associated with Inphi Corp as CTO for Storage & Computing and with Samsung Electronics as Sr. Director working on fault-tolerant DRAM memory and storage array architectures. During 2006 through 2011, Nirmal held roles as a Principal Architect, Chief Server Hardware Architect & VP at NVIDIA. From 1991 through 2009, he was also associated with Stanford University's Center for Reliable Computing and EE Department as Associate Director and Consulting Professor respectively. During his association with Stanford University, he taught courses in Logic Design, Computer Architecture, Fault-Tolerant Computing, supervised six PhD students and was co-investigator with Professor Edward J. McCluskey on DARPA's ROAR (Reliability Obtained through Adaptive Reconfiguration) project. Nirmal was the Executive VP, CTO, and COO at Alliance Semiconductor, Santa Clara. Prior to Alliance, Nirmal was VP of Architecture at Chip Engines. Nirmal has served in senior technical and management positions at Tiara Networks, Silicon Graphics, HaL Computers, and Hewlett Packard.

Nirmal received his BE ECE degree (1982) from Osmania University, India; MSEE degree (1984) from the University of Iowa; and Ph.D. EE degree (1991) from Stanford University.

He is a Fellow of IEEE (2002) and was cited for his contributions to reliable computing.