

SORS: Resilience Design Patterns: A Structured Approach to Resilience at Extreme Scale

Objectives

Speaker: Saurabh Hukerikar, Research Associate, Computer Science Research group, Oak Ridge National Laboratory (ORNL), USA

Abstract: Reliability is a serious concern for future extreme-scale high-performance computing (HPC) systems. Projections based on the current generation of HPC systems and technology roadmaps suggest the prevalence of very high fault rates in future systems. While the HPC community has developed various resilience solutions, application-level techniques as well as system-based solutions, the solution space of HPC resilience techniques remains fragmented. There are no formal methods and metrics to investigate and evaluate resilience holistically in HPC systems that consider impact scope, handling coverage, and performance & power efficiency. Few of the current approaches are portable to newer architectures and software environments that will be deployed on future systems.

In this talk, I will present a structured approach to the management of HPC resilience using the concept of resilience-based design patterns. A design pattern is a general repeatable solution to a commonly occurring problem. We identify the commonly occurring problems and solutions used to deal with faults, errors and failures in HPC systems. Each established solution is described in the form of a pattern that addresses concrete problems. We have developed a complete catalog of resilience design patterns, which provides designers with a collection of such reusable design elements. We have also defined a framework that enhances a designer's understanding of the important constraints and opportunities for the design patterns to be implemented and deployed at various layers of the system stack. This design framework may be used to establish mechanisms and interfaces to coordinate flexible fault management across hardware and software components. The framework also enables optimization of the cost-benefit trade-offs among performance, resilience, and power consumption. The overall goal of this work is to enable a systematic methodology for the design and evaluation of resilience technologies in extreme-scale HPC systems.



at the Oak Ridge National Laboratory (ORNL) in the Computer Science Research group. He received his PhD from the Department of Electrical Engineering at the University of Southern California. At USC, he worked with the Computational Systems and Technology Division at the Information Sciences Institute. His research is driven by resilience, performance, power and energy efficiency challenges for extreme-scale high-performance computing systems, including system architecture, compiler techniques, programming models and runtime system design, hardware/software co-design, performance modeling and simulation. His recent projects include the study of design patterns for HPC resilience and the offline and online analysis of failures in ORNL's Titan supercomputer.

Barcelona Supercomputing Center - Centro Nacional de Supercomputación

Source URL (retrieved on 25 abr 2024 - 09:01): <https://www.bsc.es/ca/research-and-development/research-seminars/sors-resilience-design-patterns-structured-approach-resilience-extreme-scale>