Tutorial I

Autonomic Grid Computing

Abstract

The emergence of computational Grids and the potential for seamless aggregation, integration and interactions has made it possible to conceive a new generation of realistic, scientific and engineering simulations of complex physical phenomena. These applications will symbiotically and opportunistically combine computations, experiments, observations, and real-time data, and will provide important insights into complex systems. However, the phenomenon being modeled by these applications is inherently multi-phased, dynamic and heterogeneous (in time, space, and state) requiring very large numbers of software components and very dynamic compositions and interactions between these components. Furthermore, the underlying Grid infrastructure is similarly heterogeneous and dynamic. The combination of the two results in application development, configuration and management complexities that break current paradigms based on passive components and static compositions. Autonomic computing offers a potential solution. It is inspired by nature and biological systems (such as the autonomic nervous system) that have evolved to cope with the challenges of scale, complexity, heterogeneity and unpredictability by being decentralized, context aware, adaptive and resilient. This new era of computing driven by the convergence of biological and digital computing systems and is characterized by being self-defining, self-configuring, self-optimizing, self-protecting, self-healing, context aware and anticipatory.

The overall goal of this tutorial is to introduce Autonomic Grid Computing. In this tutorial we will outline the defining research issues, present the opportunities and challenges of Autonomic Computing. We will also review the current landscape of Autonomic Computing and present case studies of autonomic systems, applications and application development and execution environments.

This tutorial is primarily targeted towards researcher, practitioners, educators and students from academia and industry in the area of parallel and distributed computing. The tutorial will present material at all levels, including introductory and overview materials for beginners or novice readers, as well as in-depth research material for more advanced attendees.

Speakers

Manish Parashar
Department of Electrical & Computer Engineering
Rutgers University
parashar@caip.rutgers.edu

Salim Hariri
Department of Electrical and Computer Engineering
The University of Arizona
hariri@ece.arizona.edu

Tutorial II

Research Challenges of Autonomic Computing: An Industry Perspective

Abstract

In a keynote address to the National Academy of Engineers at Harvard University in October 2001, IBM's senior vice president of research Paul Horn stated:

"... we face a problem springing from the very core of our success — and too few of us are focused on solving it. More than any other I/T problem, this one — if it remains unsolved — will actually prevent us from moving to the next era of computing. The obstacle is complexity ... Dealing with it is the single most important challenge facing the I/T industry."

Widespread industry recognition of this challenge is exemplified by initiatives by major players like HP (Adaptive Enterprise), IBM (Autonomic Computing), Intel (Proactive Computing), Microsoft (Dynamic Systems Initiative), and Sun (N1).

As part of its response to this challenge, IBM has developed an Autonomic Computing research program that has, in the 2 ½ years since Paul Horn's address, grown to encompass well over 100 researchers working on dozens of projects taking place in over half a dozen research labs around the world. From the beginning, customer scenarios have been invaluable in driving individual autonomic computing research projects and maintaining coherency across them.

We begin the tutorial by presenting a set of industry scenarios that illustrate the complexities inherent in managing today's I/T systems and motivate the need to make them more self-managing. Following an overview of IBM's Autonomic Computing Research program, we introduce you to several of IBM's leading researchers, who will give you a deeper look inside their individual projects in areas ranging from self-optimization to self-healing to self-configuration to human-computer interaction, several of them relying on tools and techniques from artificial intelligence and mathematics. We describe a publicly available toolkit from which you may construct autonomic components, and show you how to use Web Services to build self-managing components and systems. Finally, we discuss major research challenges that face the research community in the years ahead — challenges that will require collaboration among the best minds of academia, IBM, and other premier industry research laboratories.

The tutorial is appropriate for all attendees of ICAC '04. No special prior knowledge or experience is assumed or required. It may be of particular interest to professors or students who wish to meet and explore collaborations with IBM's leading researchers in the field of autonomic computing.

Chair

Jeffrey O. Kephart
Mgr., Agents and Emergent Phenomena
IBM Thomas J Watson Research Center
kephart@us.ibm.com