

## Title: Conquering Complexity through Managed Evolution

### Abstract:

A defining characteristic of the information age is our reliance on vast, complex, and intertwined information technology (IT) plants—consisting of large numbers of heterogeneous, high-capacity computing, communication, and storage systems as well as wide-range of software components. The rapid increase in the complexity of such IT plants, however, is fast approaching a barrier. The presence of this barrier is evident in: (1) the nearly universal reluctance of system administrators to modify existing functional systems, even when upgrades are required; (2) the widespread recognition that IT systems are brittle and respond to changes in unpredictable ways; and (3) the fact that IT plants have become very labor intensive and costly to maintain, with system administration and upgrade accounting for 60-80% of the TCO in many IT plants. Consequently, systems that we depend on to promote change are fast becoming obstacles to change.

We argue that continuous *evolution* is a key contributor to complexity. Evolution of IT plants is governed by several factors, including the need to enhance the plant to accommodate additional software functionality or users; as well as the need to meet the demands of changing requirements, changing technology and/or changing operating conditions (workload, faults, etc.). Today, evolving IT plants in a timely manner while maintaining desired levels of performance, stability, and security is an art—system evolution tasks are intuition-based and manual. The inherent complexity of hardware and software components when coupled with continuous evolution of plants often lead to scenarios where technicians understand components, but can neither predict nor control the IT plant.

The objective of the *Systems Research Lab (SRL)*, a new R&D entity being established at TRDDC, is to conquer complexity by developing methodologies and tools to *manage system evolution*. As a first step, we have defined the *Data-intensive Computing initiative* to understand and model *operational characteristics* of IT plants, and thereby simplify many tasks such as capacity planning, fault detection and isolation, anomaly detection, and performance engineering. To achieve this goal, we are developing a *unified, extensible, and customizable architecture*—referred to as the *Information Integration and Inference (I<sup>3</sup>) architecture*—for monitoring, assimilating, and analyzing large volumes of data from a diverse set of distributed and heterogeneous sources.

In this talk, I will describe the overall R&D vision for the Systems Research Lab, and then outline some of the specific activities and challenges involved in the Data-intensive Computing initiative.

### Bio:

Harrick Vin is a Vice President at the Tata Consultancy Services (TCS) and the head of the Systems Research Lab (SRL) being started at the Tata Research, Development and Design Center (TRDDC) in Pune, India. TRDDC is an R&D division of TCS. Harrick is also a Professor of Computer Sciences at the University of Texas at Austin.

Harrick's research interests are in the areas of networks, operating systems, distributed systems, and multimedia systems. Harrick received his Ph.D. in Computer Science from the University of California at San Diego in 1993. He has co-authored more than 100 papers in leading journals and conferences. Harrick is a recipient of several awards including the Faculty Fellow in Computer Sciences, Dean's Fellowship, National Science Foundation CAREER award, IBM Faculty Development Award, Fellow of the IBM Austin Center for Advanced Studies, AT&T Foundation Award, National Science Foundation Research Initiation Award, IBM Doctoral Fellowship, NCR Innovation Award, and San Diego Supercomputer Center Creative Computing Award. He has served on the Editorial Board of ACM/Springer Multimedia Systems Journal, IEEE Transactions on Multimedia, and IEEE Multimedia. He has been a guest editor for IEEE Network. He has served as the program chair, the program co-chair, and a program committee member for several conferences.