Computer Science Lecture Series

Achieving Scale and Reliability in Distributed Systems: A Probabilistic Protocol Methodology

Mr. Indranil Gupta

Cornell University

12:30-1:30pm
Wednesday, April 9, 2003
Room 381 - Computer Science Dept.

Abstract:

For several years, peer-to-peer technology has been at the focus of increasing attention among users and researchers. It is used in the design of file sharing systems (e.g., Napster, Kazaa), data centers and distributed file systems. It is likely to influence evolution of intelligent, autonomous and large-scale systems of small devices (e.g., sensor nodes). The success of such systems depends on their ability to sustain performance under dynamic stresses (network congestion, end point failures, system churn) as well as to scale gracefully when end points number into thousands or millions. The traditional approaches to designing distributed protocols (e.g., two phase commit, heartbeating) have difficulty scaling up, and are restricted in some settings to a few hundred participants and very light ranges of perturbation. I will discuss an alternative methodology to design large-scale distributed protocols. This methodology generates probabilistic protocols for a suite of problems central to the design of large-scale peer-to-peer systems. We have designed and experimented with protocols for group membership, resource location, variants of reliable multicast and consensus, and data aggregation. Probabilistic protocols impose low overheads on the group (costs at participants are often constant), and have low probabilities of incorrectness. They can be backed up with inexpensive repair protocols that ensure deterministic reliability. The scalability, reliability, simplicity, and soft real-time traits of these protocols are a good match with the hardware and application requirements in emerging areas of peer-to-peer computing such as large-scale networks of devices.