Guest Editorial: Management of Mobility in Distributed Systems

Adarshpal S. Sethi [a] Metin Feridun [b]

[a] *Department of Computer and Information Sciences*
University of Delaware, Newark, DE, USA
E-mail: sethi@cis.udel.edu
[b] *IBM Research Division, Zurich Research Laboratories*
Rueschlikon, Switzerland
E-mail: fer@zurich.ibm.com

In response to the emergence of the Internet and enterprise Intranets as the enabler for network based businesses and services, support for mobility is increasingly demanded for networks as a means for expanding the scope and reach of network-based applications. Mobile or nomadic users may now connect to an application via a variety of different channels, such as high-speed local area networks, dial up modems, mobile telephones etc. Users may connect from different stationary locations, such as an office, home, or a hotel, or may actually connect while on the move, such as in a train or a car. As a result, the distinctions between wireless and terrestrial, as well as fixed and mobile are getting blurred as the new network-based distributed systems encompass all of these within their scope. New infrastructure and services are needed to support seamless interaction between these widely varying types of network channels with widely varying quality of service in a manner that is transparent to the users of these services. Intelligent filtering and caching mechanisms are needed to adjust traffic to available bandwidth. Location services are needed to track users for location-aware applications or to redirect traffic to the nearest network access point. Management of this mobility poses significant problems that form the subject of this special issue.

The problem of Location Management has received considerable attention in the literature on mobility. This problem refers to the techniques used to keep track of the location of a mobile user so that both incoming and outgoing

communication can be efficiently routed. In this context, two different procedures are commonly used, one called Location Update in which the mobile user informs the network agents (which maintain the location databases) of its new location, and another called Location Search when the agents search for the user in order to establish an incoming call to it. Each procedure has an associated cost and many Location Management algorithms try to tradeoff the cost of one with the other in order to achieve a given cost/performance objective.

The first paper in this issue, "Agent-Based Forwarding Strategies for Reducing Location Management Cost in Mobile Networks" by Ing-Ray Chen, Tsong-Min Chen, and Chiang Lee proposes new lower-cost strategies for location management in mobile networks. These schemes are based on the concept of a local agent in order to reduce the twin costs of location management, namely those of location updates and of search, particularly for the cases where there is a high degree of mobility. The performance of the schemes is demonstrated through analytic modeling.

A slightly different approach to the problem of location management is taken by the second paper, "Optimal Distributed Location Management in Mobile Networks" by Govind Krishnamurthi, Murat Azizoğlu, and Arun K. Somani. The paper considers the class of strategies which are load balanced in the sense that the number of location updates and queries received by an information database per unit time must be equal to those of other databases on the average. This requirement avoids performance bottlenecks in the system. Krishnamurthi et al obtain lower bounds to the worst-case and average delay in locating a mobile user, and to the call blocking probability for this class of load-balanced strategies. They then propose a dynamic location management algorithm that meets all these bounds.

The third paper, "An Efficient Location Management Protocol for Wireless ATM Customer Premises Networks" by Evangelos Zervas, Alexandros Kaloxylos, and Lazaros Merakos extends mobile PNNI (Private Network-to-Network Interface) to provide routing functionality in wireless ATM networks. Although the routing protocol can be considered to duplicate the task of location management at a cost of increased computation and bandwidth, this paper proposes a modified PNNI mechanism that includes some characteristics of location management schemes in order to improve its performance and to make it more scaleable.

The next paper, "Flexible Network Support for Mobile Hosts" by Xinhua Zhao, Claude Castelluccia, and Mary Baker tackles the architectural and protocol

issues involved in providing connectivity and flexible network support for portable computers and mobile hosts as they move around a wireless network. While Mobile IP is an important first step towards supporting mobility, it does not provide the finer per-flow control of traffic that is needed for seamless, transparent mobility. This paper introduces mechanisms for supporting multiple packet delivery methods and for using multiple network interfaces simultaneously in a manner that adapts to the characteristics of each traffic flow.

The paper "Scalable Multicasting: The Core-Assisted Mesh Protocol" by Ewerton L. Madruga and J.J. Garcia-Luna-Aceves tackles the problem of constructing and maintaining the routing infrastructure for multicast traffic in the presence of mobility. The traditional approach of establishing a routing tree for multicast traffic can be inefficient if the topology of a wireless network changes dynamically due to router mobility. A better approach is to build routing graphs called multicast meshes, usually established by using flooding protocols. This paper describes and analyzes the operation of a protocol called CAMP (Core-Assisted Mesh Protocol) which builds these meshes without flooding. The results of simulation experiments show that CAMP gives better performance than tree-based multicast routing protocols for mobile networks.

The paper "Java-Based Mobile Asset Location" by Luca Deri considers the problem of locating assets such as computing resources in a dynamically changing environment that includes moving assets. Deri describes the design of an asset location system written in Java and based on the SNMP network management protocol that can determine both where a given asset is physically located and what assets are contained in a given location. The system includes a Web user interface for ease of use and performs completely automatic asset tracking in order to avoid dependency on human operators for updating asset information.

While Mobile IP and Location Management techniques aim to hide the impact of mobility on applications, the paper "Meta-Level Architectures for Component-Based Mobile Computing" by Arnel I. Periquet and Eric C. Lin describes an approach to mobile-aware application development. This approach integrates mobile-awareness into existing distributed component designs based on a set of meta-level architectures for mobility. Mobile components are proposed to provide mechanisms for migrating application services by allowing internal constituent components to be dynamically swapped at run-time.

The last paper "Indexing Techniques for Power Management in Multi-Attribute Data Broadcast" by Qinglong Hu, Wang-Chien Lee, and Dik Lun Lee

investigates the problem of managing data broadcasts on wireless channels with the objective of minimizing the power consumption of mobile clients. The paper describes schemes that support queries for data items with multiple attributes, a more realistic situation for most business applications than the single attribute data items handled by existing techniques.

The guest editors would like to thank the authors for their contributions, and the many reviewers for providing us with timely and diligent reviews. We also wish to express our gratitude to Dr. Imrich Chlamtac, Editor-in-Chief, for giving us the opportunity to put this special issue together.