Bulletin Editors

I am very proud that issue after issue of the Data Engineering Bulletin exemplifies such consistently high quality. This pride, to be fully honest, is one level indirect. It is the issue editors who select the topic, solicit the articles, and ride herd on the authors so as to get the Bulletin published. So the great bulk of the credit over the years is rightly assigned to our associate editors. My pride is in having chosen the editors so successfully.

Associate editors serve for two years. We are at a junture now in which old editors are leaving and new editors are arriving. So it is time now to thank Umesh Dayal, Johannes Gehrke, Christian Jensen, and Renée Miller for the outstanding jobs that they have done. It has been a pleasure working with them to publish the fine issues of the last two years.

It is also time to welcome new editors. There are four new editors, whose names now appear on the inside front cover: Gustavo Alonso of ETH, Minos Garofalakis of Bell Labs, Meral Ozsoyoglu of Case Western Reserve, and Jignesh Patel of Michigan. This new crop of editors continues my ”lucky streak” of succeeding in attracting outstanding editors for the Bulletin. This is the most important part of my job as editor-in-chief. The great new issues from these editors begin with the current issue, which is described below.

The Current Issue

Performance of our systems has always been important. But increasingly, it is scalability that is the ”performance” goal for systems. Implementors are willing to trade some system efficiency for an ability to scale the system to very high performance via adding additional resources. Here is where caching and replication come into play.

Maintaining data replicas consumes system execution cycles. However, replication of data enables system performance to scale by the addition of more processors and disks. The current issue explores replication and caching strategies. There are several important considerations in this. One isse is the extent to which applications need to know where data is stored or replicated. This is transparency. Another issue is one of consistency. If one accesses data stored across multiple replicas, how can one specify what is wanted in the way of consistency? Is transactional consistency needed or some other measure. And how does a system provide this consistency? Yet another issue is currency. How are replicas maintained, and how can they guarantee something with respect to how recent the information is in a replica? And how is it decided where the replicas should go?

The current issue explores the above issues and others. It brings together a mix of papers from industrial, academic and government researchers. I have long thought that having industrial participation is an important part of the Bulletin mission, even when the industrial people are *researchers*. And this is surely successful in the current issue. I am sure that readers will be well rewarded in reading about work which is actively being pursued across such a wide spectrum of institutions. Many thanks to Gustavo Alonso for bringing us this highly useful and interesting collection of articles on a very timely subject.