Changing Bulletin Editors

The Bulletin practice is to appoint editors for two years. Each editor is responsible for producing two issues, one per year, during that time. The last of the four editors of the previous two year cycle have now completed their issues and are due to ”retire”. I want to thank Amr El Abaddi and Elke Rundensteiner for doing a very capable job and producing high quality issues for Bulletin readers. I’d like to be sure that all readers are aware that each issue is, in fact, a substantial undertaking. We owe a debt to Amr and Elke for their fine issues and their hard work.

To complete the appointment of new editors for the next two years, I am delighted to announce that Luis Gravano and Gerhard Weikum have accepted my invitation to be Bulletin editors.

Luis Gravano is a faculty member in the computer science department at Columbia University. Luis has done some really inovative work in query processing and data mining, including querying over the web, and across heterogenous data collections, involving multimedia and textual as well as tabular data. Luis received his Ph.d from Stanford.

Gerhard Weikum is on the faculty at the University of the Saarlands in Germany. Gerhard received his doctoral degree from ETH in Zurich and worked for a number of years at MCC in Austin. He is currently both a TODS and a VLDB Journal editor. Gerhard has worked on a wide range of research topics, transactions, where he invented multi-level transactions, query processing, performance analysis, and recovery, on which he and I have collaborated.

About the Current Issue

Query processing is a subject that is as old as our database research community. But it is an area which continues to stir intense interest. One fundamental problem is that we never have enough information at the time that query optimization is done to do the best possible job. This has led, over a number of years now, to exploration of ”dynamic” or ”adaptive” query processing, in which the query plan attempts to take into consideration information discovered in the process of executing the query.

The editor for this issue is Alon Levy, who has assembled articles with a span from current industrial practice to cutting-edge research. The query setting varies from the traditional to heterogeneous data sources to web centric ones, from theory to architecture to operational system to reflections upon prior work. This provides readers with a sense of where the field is now, where it is going, where it has been, and why. I want to thank Alon for his very successful effort in bringing this issue to fruition.