IEEE CS Governance

Beginning in the June Bulletin, and continuing with the September issue, I have been keeping you abreast of a proposed change in how a part of the IEEE Computer Society governs itself. In particular, a governance subcommittee of the Technical Activities Committee (TAC) was formed to propose changes in how its chair is selected. All chairs of Technical Committees (TCs) are members of this committee, and I, as chair of the TC on Data Engineering thus attend its meetings. Under the current rules, the chair is appointed, but under the proposal from the governance committee, the chair would be elected by TAC members.

I am happy to report that the TAC members have now voted to approve this proposal. So now the proposal goes forward to the Computer Society Board of Governors. We are hopeful that they will approve, and if so, the proposal becomes the new way of choosing the TAC chair. I belonged to the subcommittee and am a firm believer in the proposed change, as it gives TCs a way to have their views more effectively expressed, both in our TAC meeting and because the TAC chair represents the TCs at higher levels of the Computer Society. So we are not done yet, but progress has been made. Stay tuned.

The Current Issue

Database systems have been notoriously cranky to deal with, requiring skilled professionals to manage and keep tuned. The research into self-tuning database systems, while not eliminating all such manual involvement, reduced the level of effort considerably. This has permitted the classic database market to service much smaller enterprises than it had earlier. And, of course, it made life simpler for many more enterprises, large and small.

This "self tuning" research has been going on now for over ten years. The result has been that the "friendliness" of databases has improved greatly. But it is still the case that databases can be a challenge to deal with. So interest in the field has remained high, leading to the establishment, a few years back, of the Data Engineering Workgroup on Self-Managing Database Systems within the Technical Committee on Data Engineering.

The current issue has been assembled by a pair of editors from the workgroup, Shivnath Babu and Kai-Uwe Sattler. They know well the self-managing databases research community, the researchers, the practitioniers, and their work, and they have successfully brought together a blend of papers that reflect this diversity, with a heavier dose of industrial participation than usual. This is a reflection that this topic has commercial importance to the database industry, and hence is the kind of issue that might only come from the Data Engineering Bulletin. I want to thank Shivnath and Kai-Uwe for their efforts in bringing this fine issue to fruition.