We don't sort twice as fast every 18 months on the same computer. However, processors have on average improved at about this rate for some four decades—a phenomenon known as Moore's Law, although it is not a law, but is better characterized as a self-fulfilling prophecy. Among other kinds of hardware close to the heart of a database researcher, disks and networks improve at even faster rates than do processors. These and many other advances in hardware are important drivers for research in software, including data management research.

Highly portable and unobtrusive, wirelessly on-line electronics with sensing capabilities are appearing on the horizon. This has brought increasing interest to areas such as ubiquitous and pervasive computing, mobile computing, assistive computing, spatio-temporal databases, sensor-data management, stream data processing, augmented reality, and ambient intelligence.

In step with applications and services being delivered increasingly to their users via mobile computing devices, instead of via desk-top computers, new kinds of services with new characteristics become of interest. Context awareness is particularly important to mobile users who find themselves in a range of different situations, characterized by diverse, specific needs. Location, the capture of which is made possible through increasingly sophisticated positioning technologies, is an essential aspect of context awareness.

The specific focus of the present issue of the Data Engineering Bulletin is on a range of aspects of the indexing of the positions of continuously moving objects.

The positions of moving objects are obtained via some form of sampling, and the past positions of a point object moving in d-dimensional space are frequently represented as a polyline, a sequence of line segments in $d + 1$-dimensional space connecting the time-referenced, sampled positions. In the first paper, Dieter Pfoser considers the problem of indexing such trajectories. Focusing also on information relating to the past, Papadias et al. in the next paper consider indexing in the context of aggregate computation for moving objects.

A somewhat different problem is that of indexing the positions of moving objects from the times their positions were last sampled and into the future. Here, it is typical to represent the positions of an object by a linear function. Papadopoulos et al. study the use of duality transforms for the indexing of such positions of objects moving in one-, so-called 1.5-, and two-dimensional spaces. Agarwal and Procopiuc cover both problems, surveying moving-object indexing from a computational geometry perspective and thus broadening this issue to also cover highly relevant results not published in typical database outlets. Šaltenis and Jensen point the attention to the need for faster update processing in indices for the current positions of moving objects.

Chon et al. proceed to consider the management of moving-object trajectories. They adopt a partitioning-based technique in place of indexing. Next, when attempting to develop a practical index with high query and update performance, empirical performance studies are of essence. Myllymaki and Kaufman describe a testbed for dynamic spatial indexing, an important piece of infrastructure in this regard. Finally, Kothuri and Ravada explore the support in the Oracle DBMS for spatio-temporal indexing.

This issue samples research results and also points to challenges in an area with many open problems. It is my hope that the issue offers a good feel for the breadth of fundamental problems and challenges inherent to moving-object indexing, and for what it entails to conduct research in this area. I also hope that the issue will inspire new research on moving-object indexing.