Data replication is one of the oldest and most persistent topics of research in distributed databases. It holds the promise of solving both problems of fault-tolerance and performance. By replicating the data on remote sites, failures may occur and users may still be able to access data; alternatively, if a data item is replicated on a local or nearby site, accesses to that copy are efficient and may involve very low overhead. To obtain these benefits, however, complex and expensive synchronization mechanisms are often needed to maintain the consistency and integrity of data. In fact it has often been argued that due to the synchronization overheads and the increased possibility of deadlock, replication is impractical. Nevertheless, the appeal of replication persists and sometimes even migrates to commercial industrial products. In this special issue various researchers and practitioners will address the current state of replication and its promises.

The first paper, written by Avishai Wool, addresses head-on the promise of replication and asks the question: Is it science or fiction? He explores this question in the context of quorum-based replication and discusses its potential to solve the problems of managing replicated databases. Avishai explains why quorum systems have not fulfilled their old promises, but at the same time argues why new technological advances and applications have brought new requirements that can be addressed by quorums. In fact, he shows that quorum systems may offer a way to scale up throughput in heavily loaded systems.

Many replicated storage systems have advocated relaxed consistency models to get improved availability, scalability and performance. Doug Terry, Karin Petersen, Mike Spreitzer and Marvin Theimer discuss the case for non-transparent replication using the Bayou system. The Bayou system developed at Xerox PARC is an example of a replicated storage system that was designed to strike a balance between application control and complexity. In this paper the authors discuss the tension between overburdening and underpowering applications. They discuss two bayou applications, a calendar manager and a mail reader, and illustrate ways in which they utilize Bayou's features to manage their data in an application-specific manner.

The web is a natural context for replication with its increasing scalability and performance demands. Micha Rabinovich examines various issues that arise in Web content replication, paying special attention to challenges in dynamic replication. The paper starts with an overview of the current state of replication and caching on the web. It then proceeds to offer valuable insights into open questions and research issues both from architectural as well as from algorithmic points of view.

In the next paper, Michel Raynal addresses the problem of atomic commitment in distributed replicated databases. He shows that solutions to the consensus problem can be very useful in developing solutions to the atomic commitment problem. In particular, he develops a solid theoretical foundation for identifying the minimal assumptions a system must satisfy in order to develop non-blocking replica control protocols in spite of process crashes and asynchrony.

Finally, Bob Breton presents the evolution of a replication system, Sybase's Replication Server, and discusses how it supports disaster recovery as well as support for high availability business requirements. In particular, he discusses the need for replication as a means for disaster recovery in real applications that have very stringent availability requirements. In the paper he also discusses a number of alternatives for high availability database architectures and outlines the main risks and benefits of each approach.