

Abstract:

The ability to store data persistently is one of the key elements of information technology. In modern computers this task is usually realised by magnetic hard disk drives. They form the last and, therefore, the slowest level in the memory hierarchy. Unfortunately, single disks can not store the annually growing data volumes. Therefore, this thesis introduces the concept of storage networks, a collection of disks connected by an arbitrary network. Storage networks have a lot of advantages like parallel access to data and almost unlimited capacity growth but they can only be explored by an adaptive data distribution. This thesis introduces and analyses balancing and adaptive data distribution strategies and proves their quality theoretically. In the second part, these strategies are applied to the storage virtualisation problem, i.e. the decoupling of the logical view from the physical representation of a storage network. The resulting prototype, implemented mostly as a block device driver under LINUX, is finally evaluated in real-world scenarios.