OTTO VON GUERICKE UNIVERSITÄT MAGDEBURG

Parallel Computing and I/O

A Common Storage Engine for Modern Memory and Storage Hierarchies (SMASH)

Scientific research is increasingly driven by data-intensive problems. As the complexity of studied problems is rising, so does the need for high data throughput and capacity. The globally produced data volume doubles approximately every two years, leading an exponential data deluge. This deluge then directly challenges database management systems and file systems, which provid the foundation for efficient data analysis and management. These systems use different memory and storage devices, which we traditionally divided into primary, secondary and tertiary memory. However, with the introduction of the disruptive technology of non-volatile RAM (NVRAM), these classes started to merge into one another leading to heterogeneous storage architectures, where each storage device has highly different performance characteristics (e.g., persistence, storage capacity, latency). Hence major challenge is how to exploit the specific characteristics of memory devices.

To this end, SMASH will investigate the benefits of a common storage engine that manages a heterogeneous storage landscape including traditional storage devices and non-volatile memory technologies. The core for this storage engine will be B-epsilon-trees, as they can be used to efficiently exploit these different devices. Furthermore, data placement and migration strategies wil be investigated to minimize the overhead caused by transferring data between different devices. Eliminating the need for volatile caches will allow data consistency guarantees to be improved. From the application side, the storage engine will offer key-value and object interfaces that can be used for a wide range of use cases, such as high-performance computing (HPC) and database management systems. Moreover, due to the widening gap between the performance of computing and storage devices as well a their stagnating access performance, data reduction techniques are in high demand to reduce the bandwidth requirements wher storing and retrieving data. We will, therefore, conduct research regarding data transformations in general and the possibilities o external and accelerated transformations. As part of SMASH, we will provide a prototypical standalone software library to be use by third-party projects. Common HPC workflows will be supported through an integration of SMASH into the existing JULEA storage framework, while database systems can use the interface of SMASH directly whenever data is stored or accessed.

This project is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 502268500.

- Prof. Dr. Gunter Saake
- Dr. David Broneske

- https://gepris.dfg.de/gepris/projekt/502268500
- https://github.com/julea-io