論文の内容の要旨
Abstract

論文題目　A Decentralized Implementation of Software Distributed Shared Memory（ソフトウェア分散共有メモリの非集中型実装）

氏名　遠藤 亘

The low application productivity of supercomputers has been a remaining problem in the history of parallel computing. Most of today's supercomputers are distributed-memory machines, which force application programmers to manually issue communication requests. Shared memory, the other memory model, is considered more productive than distributed memory, but it is widely believed that shared-memory programming is not scalable due to the model itself. In the 1990s, many researchers had developed software systems called Distributed Shared Memory (DSM) systems that realize the shared-memory model on top of distributed-memory machines. Nowadays, however, only very few researchers are investigating this idea because of the observed results of its poor performance.

The hardware environment for supercomputing has been largely improved for two or three decades. The interconnection network between nodes has also been significantly improved. The technology trends of both hardware and software have enlightened the possibility of better DSM systems than traditional DSM systems. For example, there were very few DSM systems that utilized the interconnection feature called Remote Direct Memory Access (RDMA), which appeared in the 2000s. Most of RDMA implementations can only support point-to-point communications and it is necessary to implement a decentralized coherence protocol for DSMs to utilize RDMA. Multi-core architectures have become dominant in today's processors and utilizing multi-threading in a node possibly improves the performance of DSM systems, but traditional DSM systems are not optimized for modern hardware.

In this dissertation, we propose a decentralized approach to implement distributed shared memory based on RDMA and multi-core architecture. This approach is demonstrated as a software DSM system Menps. To decentralize the coherence protocol of Menps, we have introduced three novel ideas: the floating home-based method, the hybrid invalidation of logical timestamps and write notices, and the fast release. The evaluation results of Menps show that the proposed coherence protocol can improve the performance of software distributed shared memory compared with the existing methods.

To exploit the performance of RDMA on multi-core architectures, we have developed a new communication library MECOM for Menps. MECOM is a communication library based on a new technique

to implement software communication offloading using user-level threading. The proposed software offloading technique using atomic operations improves the message rates of RDMA communications in multi-threading environments compared with the conventional locking schemes.

To accelerate the thread scheduling in MENPS, we also have implemented a user-level threading library ComposableThreads. ComposableThreads provides compile-time parametricity for different purposes in system programming. This library can be not only employed as the tasking layer for MENPS, but also customized as an efficient general-purpose threading library.

This dissertation describes the design and implementation of the DSM library MENPS. As an integrated research prototype for transparently executing shared-memory multi-threading applications, MENPS provides many interesting insights about how we can design parallel computer architecture.