Letter from the Special Issue Editor

High-performance Remote Direct Memory Access (RDMA) capable networks such as Infiniband FDR/EDR are fundamentally changing the design of distributed data centric systems. Until recently, such systems were built on the assumption that the network is the main bottleneck. Thus, the systems aim to reduce communication between nodes using techniques such as locality-aware partitioning schemes, semi-joins, and complicated preprocessing steps. Yet, with next generation technologies, the network is no longer the dominant bottleneck.

Even today, bandwidth from InfiniBand FDR 4× is roughly that of one memory channel. DDR3 memory bandwidth ranges from 6.25GB/s (DDR3-800) to 16.6GB/s (DDR3-2133) per channel, whereas InfiniBand has a bandwidth of 1.7GB/s (FDR 1×) to 37.5GB/s (EDR 12×) per NIC port (Figure 1). Future InfiniBand standards (e.g., HDR, NDR) promise to dramatically exceed memory channel bandwidth. Furthermore, while modern CPUs often take advantage of several memory channels (usually 4 per socket), dual-port NICs are becoming the standard and in contrast to memory channels, these NICs support full duplex rather than half-duplex. Thus, today a single dual-port InfiniBand EDR 4× NIC roughly matches the bandwidth of the standard 4 memory channels of a single CPU socket with a read/write workload.

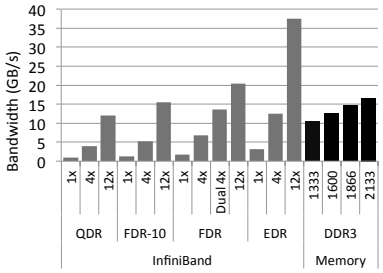

Figure 1: Memory vs Network Bandwidth

RDMA advances have also improved network latency significantly. Our InfiniBand FDR 4x experiments show that it takes only ≈ 1 micro second to transfer 1KB of data using RDMA, compared ≈ 0.08 micro seconds required by the CPU to read the same amount of data from main memory (RAM). With only 256KB, there is virtually no difference between the access times since the bandwidth starts to dominate the transfer time. Nonetheless, one cannot assume that a high-performance network changes a cluster into a NUMA architecture because: (1) RDMA-based memory access patterns differ from local memory access patterns; (2) random access latency for remote requests remains significantly higher; and (3) hardware-embedded coherence that ensures data consistency in NUMA does not exist for RDMA. How RDMA and next generation networks influence data management systems is an open question and the topic of this special issue.

The first paper, **RDMA Reads: To Use or Not to Use?** by Dragojević et al discusses a fundamental question of using the next generation networks: is RDMA or RPC the right building block for remote reads. RDMA reads perform better than RPCs in many cases, but RPCs can perform additional operations before replying to the client. **Designing Databases for Future High-Performance Networks** by Barthels et al makes a case that the recent advances in network technologies enable a re-evaluation and re-design of several system design concepts and database algorithms. Furthermore, the authors argue that databases and networks should be co-designed in the future and require an extended network instruction set architecture (NISA) to better support distributed database designs. **Rethinking Distributed Query Execution on High-Speed Networks** by Salama et al reconsiders traditional query execution models for next generation network technology and presents an end-to-end solution based on the network-attached-memory (NAM) architecture. **Crail: A High-Performance I/O Architecture for Distributed Data Processing** by Stuedi et al presents an I/O architecture redesigned for highspeed networks and RDMA, especially targeting temporary data that requires great performance for operations such as data shuffling or broadcasting. Their solutions allow, for example, the Apache data processing ecosystem to fully exploit new network technology. The last two papers, **Scalable and Distributed Key-Value Store-based Data Management Using RDMA-Memcached** by Lu et al and **Beyond Simple Request Processing with RAMCloud** by Kulkarni et al are concerned with leveraging RDMA for key/value stores with very impressive latency and performance results.

I would like to thank the authors for their insightful contributions to this special issue. Happy reading!

Tim Kraska
Brown University