ferent compared to RDMA over InfiniBand. The designed *InterSense* emulator offers unique capabilities for analysis of scale-out distributed memory applications.

7. CONCLUSION AND FUTURE WORK

In this work, we introduce novel bandwidth and latency control mechanisms for performance emulation of the high-speed interconnects. We built a prototype of a new emulator and carefully evaluated its performance, efficiency, and accuracy. *InterSense* can assist researchers and engineers in emulating a variety of performance characteristics of future large-scale interconnects and conducting the application sensitivity and scalability analysis dependent on these characteristics.

We are working on augmenting the proposed approach with additional profiling, modeling, and prediction technique. By performing the emulation in small deployments with increased interconnect latency and decreased bandwidth we aim to derive the predictive models for application performance when processing larger data amounts in large-scale distributed environments. We believe that the *InterSense* ability to accurately indicate the *needed* interconnect bandwidth for achieving the user-defined application performance objectives and to reflect the application sensitivity to the increased interconnect latency will help in applications' optimization and redesign.

8. REFERENCES

[1] Emulab - Network Emulation Testbed, `http://www.emulab.net/` .

[2] Graph 500 Benchmark. `www.graph500.org/` .

[3] HP Labs. The Machine: A new kind of computer. `http://www.hpl.hp.com/research/systems-research/` .

[4] HPCC RandomAccess (GUPS) Benchmark. `http://icl.cs.utk.edu/projectsfiles/hpcc/RandomAccess/` .

[5] MVAPICH: MPI over InfiniBand, 10GigE/iWARP and RoCE. `http://mvapich.cse.ohio-state.edu/`.

[6] MVAPICH Ohio State University Micro benchmark. `http://mvapich.cse.ohio-state.edu/benchmarks/`.

[7] NAS Parallel Benchmarks. `http://www.nas.nasa.gov/publications/npb.html` .

[8] netem, `http://www.linuxfoundation.org/collaborate/workgroups/networking/netem` .

[9] K. Asanovic. FireBox: A Hardware Building Block for 2020 Warehouse-Scale Computers. In *Proc. of FAST*, 2014.

[10] G. Banga, J. C. Mogul, and P. Druschel. A scalable and Explicit Event Delivery Mechanism for UNIX. In *Proc. of the USENIX Annual Technical Conference*, 1999.

[11] F. Checconi and F. Petrini. Traversing Trillions of Edges in Real Time: Graph Exploration on Large-Scale Parallel Machines. In *Proc. of Intl. Parallel and Distributed Processing Symposium, IPDPS'14*, 2014.

[12] F. Checconi, F. Petrini, J. Willcock, A. Lumsdaine, A. R. Choudhury, and Y. Sabharwal. Breaking the Speed and Scalability Barriers for Graph Exploration on Distributed-Memory Machines. In *Proc. of Conference on High Performance Computing Networking, Storage and Analysis, SC'12,*, 2012.

[13] J. Jose, H. Subramoni, M. Luo, M. Zhang, J. Huang, M. Wasi-ur Rahman, N. S. Islam, X. Ouyang, H. Wang, S. Sur, and D. K. Panda. Memcached Design on High Performance RDMA Capable Interconnects. In *Proc. of the 2011 International Conference on Parallel Processing*, ICPP '11, 2011.

[14] X. Lu, M. Wasi-ur Rahman, N. S. Islam, D. Shankar, , and D. K. D. Panda. Accelerating Spark with RDMA for Big Data Processing: Early Experiences. In *Proc. of Hot Interconnects*, 2014.

[15] R. P. Martin, A. M. Vahdat, D. E. Culler, and T. E. Anderson. Effects of Communication Latency, Overhead, and Bandwidth in a Cluster Architecture. In *Proc. of the 24th Annual International Symposium on Computer Architecture*, ISCA '97, 1997.

[16] E. M. Nahum, M.-C. Rosu, S. Seshan, and J. Almeida. The Effects of Wide-area Conditions on WWW Server Performance. In *Proc. of the 2001 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '01, 2001.

[17] B. D. Noble, M. Satyanarayanan, G. T. Nguyen, and R. H. Katz. Trace-Based Mobile Network Emulation. In *Proc. of SIGCOMM*, 1997.

[18] X. Que, F. Checconi, and F. Petrini. Performance Analysis of Graph Algorithms on P7IH. In *Proc. of the 29th Intl. Conference on Supercomputing, ISC'14*, 2014.

[19] V. Saxena, Y. Sabharwal, and P. Bhatotia. Performance evaluation and optimization of random memory access on multicores with high productivity. In *Proc. of Intl. Conference on High Performance Computing (HiPC)*, 2010.

[20] A. Vahdat, K. Yocum, K. Walsh, P. Mahadevan, D. Kostić, J. Chase, and D. Becker. Scalability and accuracy in a large-scale network emulator. *SIGOPS Oper. Syst. Rev.*, 36(SI), Dec. 2002.

[21] Q. Wang, L. Cherkasova, J. Li, and H. Volos. InterSense: Interconnect Performance Emulator for Future Scale-out Distributed Memory Applications. In *Intl. Symposium on Modelling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*, 2015.

[22] M. Wasi-ur Rahman, N. S. Islam, X. Lu, J. Jose, H. Subramoni, H. Wang, and D. K. D. Panda. High-Performance RDMA-based Design of Hadoop MapReduce over InfiniBand. In *Proc. of the 2013 IEEE 27th International Symposium on Parallel and Distributed Processing Workshops and PhD Forum*, IPDPSW '13, 2013.

[23] M. Wasi-ur-Rahman, X. Lu, N. S. Islam, R. Rajachandrasekar, and D. K. Panda. MapReduce over Lustre: Can RDMA-Based Approach Benefit? In *Proc. of the 20th International Conference EuroPar*, 2014.

[24] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb, and A. Joglekar. An Integrated Experimental Environment for Distributed Systems and Networks. *SIGOPS Oper. Syst. Rev.*, 36(SI), Dec. 2002.

[25] H. Yu and A. Vahdat. The Costs and Limits of Availability for Replicated Services. In *Proc. of the 18th ACM Symposium on Operating Systems Principles (SOSP)*, 2001.