

we use Infiniband network, we may not experience more latency delay or throughput degradation.

The mechanism we suggest might not be cost effective for someone who frequently runs large memory application and require very fast response time. Even though it has overhead of user-level page fault handling, it may be a silver bullet for someone who has cluster system and need to run large memory application sometimes. It would be a good solution for someone who develops applications with shared memory among cluster machines.

ACKNOWLEDGMENT

This work was partly supported by the ICT R&D program of MSIP/IITP. [B0101-15-0104, The Development of Supercomputing System for the Genome Analysis] and the 'Cross-Ministry Giga KOREA Project' of the Ministry of Science, ICT and Future Planning, Republic of Korea (ROK). [GK14P0100, Development of Tele-Experience Service SW Platform based on Giga Media]

REFERENCES

- [1] Douglas Comer, "A New Design for Distributed Systems: The Remote Memory Model", in Proc. Of the USENIX Summer Conference. Ahaheim, California. Pp127-135, 1990
- [2] V. Roussev, G. G. Richard III, and D. Tingstrom, "dRamDisk: Efficient RAM sharing on a commodity cluster," in 25th IEEE International Performance, Computing, and Communications Conference (IPCCC 2006), Phoenix, Arizona, Apr. 10-12, 2006, pp. 193-198.
- [3] L. Iftode, K. Petersen, and K. Li, "Memory Servers for Multicomputers," in IEEE COMPCON'93 Conference, February 1993.
- [4] E. Anderson and J. Neefe. An Exploration of Network RAM. Technical Report CSD-98-1000, UC Berkley, 1998.
- [5] S. Koussih, A. Acharya, and S. Setia. Dodo: A User-Level System for Exploiting Idle Memory in Workstation Clusters. In Proceedings of the Eighth IEEE International Symposium on High Performance Distributed Computing, 1999.
- [6] McDonald, I.: Remote paging in a single address space operating system supporting quality of service. Tech. Report, Dept. of Comp. Science, Univ. of Glasgow (1999)
- [7] Y. Jeegou. Implementation of Page Management in Mome, a User-Level DSM. In Proc. of the 3rd IEEE International Symposium on Cluster Computing and the Grid (CCGRID '03), Tokyo, Japan, May 2003.
- [8] S. Liang, R. Noronha, and D. K. Panda, "Swapping to remote memory over infiniband: an approach using a high performance network block device," in IEEE Cluster Computing, 2005.
- [9] M. Dahlin, R. Wang, T. E. Anderson, and D. A. Patterson, "Cooperative caching: Using remote client memory to improve file system performance," in Operating Systems Design and Implementation, 1994.
- [10] Ming Zhao, Dynamic policy disk caching for storage networking, IBM Research Report. Publication Date: Nov 2006.
- [11] J. Liu, W. Huang, B. Abali, and D. K. Panda, "High performance VMM-bypass I/O in virtual machines," in Proceedings of the annual conference on USENIX '06 Annual Technical Conference (USENIX ATC '06), 2006.
- [12] Haogang Chen, Xiaolin Wang, Zhenlin Wang†, Xiang Wen, Xinxin Jin, Yingwei Luo, Xiaoming Li, REMOCA: Hypervisor Remote Disk Cache, in IEEE International Symposium on Parallel and Distributed Processing with Applications, pp. 162-169, 2009
- [13] J. Oleszkiewicz, L. Ziao, and Y. Liu, "Parallel network RAM: Effectively utilizing global cluster memory for large data-intensive parallel programs," in IEEE 2004 International Conference on Parallel Processing (ICPP'04), 2004.

- [14] L. Wang, J. Zhan, and W. Shi, "In Cloud, Can Scientific Communities Benefit from the Economies of Scale?" IEEE Transactions on Parallel and Distributed Systems, vol. 99, no. PrePrints, 2011.
- [15] Hyuck Han, Young Choon Lee, Member, IEEE, Woong Shin, Hyungsoo Jung, Heon Y. Yeom, Member, IEEE, and Albert Y. Zomaya Fellow, IEEE, "Cashing in on the Cache in the Cloud" in IEEE Trans. On Parallel and Distributed Systems, Vol. 23 no. 8, pp. 1387-1399, Aug. 2012
- [16] J. Ousterhout, P. Agrawal, D. Erickson, C. Kozyrakis, J. Leverich, D. Mazières, S. Mitra, A. Narayanan, G. Parulkar, M. Rosenblum, S. M. Rumble, E. Stratmann, and R. Stutsman, "The case for RAMClouds: scalable high-performance storage entirely in DRAM," SIGOPS Oper. Syst. Rev., vol. 43, pp. 92-105, January 2010.
- [17] T. Newhall, D. Amato, and A. Pshenichkin, "Reliable adaptable network ram," in Proceedings of IEEE Cluster'08, 2008.
- [18] J. Liu, J. Wu, and D. K. Panda. High Performance RDMA-Based MPI Implementation over InfiniBand. International Journal of Parallel Programming, 32(3):167-198, 2004.



Shinyoung Ahn was born in South Korea in 1974. He received the B.E., M.E. degree in information engineering from SungKyunKwan University, Seoul, Korea, in 1997, 1999, respectively. He also received the M.E. degree in software engineering from Carnegie Mellon University, Pittsburgh, USA, in 2005.

He joined Electronics and Telecommunications Research Institute(ETRI), Daejeon, Korea, in 1999. Since 1999, he has been with the cloud computing department, where he is currently a senior researcher. His main areas of research interest are high performance computing, cloud computing, workflow scheduling, and software architecture.

Mr. Ahn is a member of Korea Information Processing Society.



Gyuil Cha was born in South Korea in 1970. He received the B.S., M.S. degree in Computer Science from Korea University, Seoul, Korea, in 1998, 2000, respectively.

He joined Electronics and Telecommunications Research Institute (ETRI), Daejeon, Korea, in 2000. Since 2011, he has been with High-Performance Computing Research Section, where he is currently a senior research member of engineering staff. His main areas of research interest are High Performance Computing (HPC), System Architecture, and Kernel software.



Youngho Kim was born in South Korea in 1973. He received the B.E., M.E. degree in Information and Communication Engineering from Chungbuk National University, Korea, in 1999 and 2001 respectively.

He joined ETRI (Electronics and Telecommunications Research Institute) in 2001. Since 2001, he has been working as a senior researcher at the Cloud Computing Department. His current research interests include High Performance Computing, Cloud Computing, and Parallel and Distributed Systems.

Mr. Kim is a member of Korea Information Processing Society.



Eunji Lim received the B.E., M.E. degree in Computer Science from Pusan National University, Busan, Korea, in 1999, 2001, respectively. Since 2001, she has been with Cloud Computing Department in Electronics and Telecommunications Research Institute(ETRI), Korea, where she is currently a senior researcher. Her main areas of research interest are Distributed System and High Performance Computing