selection is based on the data items and the operations of the transactions, the data versions available, and the network characteristics of the WS farm.

We show that although our solution can be viewed as a special case of hybrid update propagation method, it provides several advantages over existing solutions.

First, our approach provides the scalability required by modern applications, such as web services, and is suitable for the architectures and technologies implementing these applications. The buddy selection algorithm supports dynamic master-slave site selection for data items, ensures correct transaction execution, and aids load-balancing among the replicas. Finally, incorporating network specific characteristics, such as distance and bandwidth, reduces latency observed by the client. Our approach has a similar message complexity compared to a traditional lazy master-slave replication using group communication and less message complexity compared to a lazy master-slave replication without group communication.

We further extend the buddy system to handle coarse-grained web services. Our solution is based on extending UML specifications with stereotypes to embed CRUD, Parallel and data element semantics into the model. The dispatcher can then extract the semantics from the model and distribute the requests to clusters as it did with the fine-grained web service. Each individual transaction is applied to a pair of clusters synchronously allowing enforcement of consistency guarantees and durability. There are two limitations of our work we plan to address in the future. First, we are relying on the back end database for constraint guarantees. This limits the granularity of transaction parallelism. Second, we require a transaction to be sent in a single request to the dispatcher. We plan to extend the buddy system to allow the dispatcher to receive WS-Atomic and WS-BusinessActivity transactional requests that can span a single request.

References

[1] S. Gilbert and N. Lynch, "Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services," *SIGACT News*, vol. 33, pp. 51-59, 2002.

[2] D. Abadi, "Consistency tradeoffs in modern distributed database system design: Cap is only part of the story," *Computer*, vol. 45, pp. 37-42, 2012.

[3] A. Fekete, S. N. Goldrei and J. P. Asenjo, "Quantifying isolation anomalies," in *Proceedings of the VLDB Endowment*, 2009.

[4] O. M. Group, "Unified Modeling Language: Supersturcture," 05 02 2007. [Online]. Available: http://www.omg.org/spec/UML/2.1.1/. [Accessed 08 01 2013].

[5] E. Christensen, F. Curbera, G. Meredith and S. Weerawarana, "Web service definition language (WSDL)," 2001. [Online]. Available: http://www.w3.org/TR/wsdl.

[6] J. Jang, A. Fekete and P. Greenfield, "Delivering promises for web," in *Web Services, IEEE International Conference on*, 2007.

[7] M. Y. Lou and C. S. Yang, "Constructing zero-loss web services," *INFOCOM*, pp. 1781-1790, 2001.

[8] M. T. Ozsu and P. Valduriez, Principles of Distributed Database Systems, 3rd ed., Springer, 2011.

[9] Y. Lin, B. Kemme, M. Patino Martiınez and R. Jimenez-Peris, "Middleware based data replication providing snapshot isolation," in *Proceedings of the 2005 ACM SIGMOD international conference on Management of data, ser. SIGMOD ’05*, New York, NY, 2005.

[10] H. Jung, H. Han, A. Fekete and U. Rhm, "Serializable snapshot isolation," *PVLDB*, pp. 783-794, 2011.

[11] Y. Breitbart and H. F. Korth, "Replication and consistency: being lazy helps sometimes," *Proceedings of the sixteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems, ser. PODS ’97,* pp. 173-184, 1997.

[12] K. Daudjee and K. Salem, "Lazy database replication with ordering," in *Data Engineering, International Conference on*, 2004.

[13] S. Jajodia and D. Mutchler, "A hybrid replica control algorithm combin-ing static and dynamic voting," *IEEE Transactions on Knowledge and Data Engineering,* vol. 1, pp. 459-469, 1989.

[14] D. Long, J. Carroll and K. Stewart, "Estimating the reliability of," *IEEE Transactions on,* vol. 38, pp. 1691-1702, 1989.

[15] L. Irun-Briz, F. Castro-Company, A. Garcia-Nevia, A. Calero-Monteagudo and F. D. Munoz-Escoi, "Lazy recovery in a hybrid database replication protocol," in *In Proc. of XII Jornadas de Concurrencia y Sistemas Distribuidos*, 2005.

[16] A. Lakshman and P. Malik, "Cassandra: a decentralized structured," *SIGOPS Oper. Syst. Rev.*, vol. 44, pp. 35-40, 2010.

[17] A. Olmsted and C. Farkas, "The cost of increased transactional correctness and durability in distributed databases," in *13th International Conference on Information Reuse and*, Los Vegas, NV, 2012.

[18] A. Olmsted and C. Farkas, "High Volume Web Service Resource Consumption," in *Internet Technology and Secured Transactions, 2012. ICITST 2012*, London, UK, 2012.