

# ZHT: a zero-hop distributed hashtable

Ekta Joshi<sup>1</sup>, Prof. Tushar J Raval<sup>2</sup>, Prof. Karishma A chaudhary<sup>3</sup>

<sup>1</sup> M.E. Student, Computer engineering, L.D. College of Engineering, Gujarat, India

<sup>2</sup> Associate Professor, Computer Engineering Department, L D College of Engineering, Gujarat, India

<sup>3</sup> Assistant Professor, Computer Engineering Department, L D College of Engineering, Gujarat, India

## ABSTRACT

Over the last decade, storage systems have experienced a 10-fold increase between their capacity and bandwidth. This gap is predicted to grow faster with exponentially growing concurrency levels, with future exascales delivering millions of nodes and billions of threads of execution. A critical component of future file systems for high-end computing is metadata management. This extended abstract presents ZHT, a zero-hop distributed hash-table, which has been tuned for the specific requirements of high-end computing. ZHT aims to be a building block for future distributed systems, such as parallel and distributed file systems, distributed job management systems, and parallel programming systems. ZHT has some important properties, such as being light-weight, dynamically allowing nodes join and leave, fault tolerant through replication, persistent, scalable, and supporting unconventional operations such as append (providing lock-free concurrent key/value modifications) in addition to insert/lookup/remove. The primary goal of ZHT is excellent availability, fault tolerance, high throughput, and low latencies.

**Keyword:** - Distributed hash tables, key/value stores, high-end computing.

## 1. INTRODUCTION

The current architecture of high-end computing (HEC) systems is decades-old and has persisted as we scaled from gigascales to petascales. In this architecture, storage is completely segregated from the compute resources and are connected via a network interconnect. This approach will not scale several orders of magnitude in terms of concurrency and throughput, and will thus prevent the move from petascales to exascale. One of the major bottlenecks in current state-of-the-art storage systems is metadata management. Metadata operations on parallel file systems can be inefficient at large scale. Early experiments on the BlueGene/P system at 16K-core scales shows the various costs (wall-clock time measured at remote processor) for file/directory create on GPFS [1]. Ideal performance would be to have a flat line. If care is not taken to avoid lock contention, performance degrades rapidly, with operations (e.g. create directory, create file, etc) that took milliseconds on a single core, taking over 1000 seconds at 16K-core scales. [3, 4].

ZHT has been tuned for the specific requirements of highend computing (e.g. trustworthy/reliable hardware, fast networks, non-existent "churn", low latencies, and scientific computing data-access patterns), and we believe it is a good abstraction to build distributed metadata management for distributed file systems.

### 1.1 Proposed work

Future storage systems for high-end computing should support distributed metadata management, leveraging distributed data-structure tailored for this environment. The distributed data-structures share some characteristics with structured distributed hash tables [24], having resilience in face of failures with high availability; however, they should support constant time inserts/lookups/removes delivering low latencies typically found in centralized metadata management. Replication is used to ensure both metadata and data availability.

We developed ZHT, a zero-hop distributed hash-table, which has been tuned for the specific requirements of high-end computing (e.g. trustworthy/reliable hardware, fast networks, non-existent "churn", low latencies, and scientific

computing data-access patterns). The primary goal of ZHT is excellent availability, fault tolerance, high throughput, and low latencies. The report also discusses the use of ZHT as the distributed metadata management in the FusionFS distributed filesystem.

## 2. RELATED WORK

DHTs have an important role in building support for scalable meta-data across extreme scale systems. Some of the DHTs from the literature are Kademia[11], CAN[12], Chord [13], Pastry[14], Tapestry[15], Memcached[16], Dynamo[2], Cycloid[17], Ketama[18], RIAK[19], Maidsafe-dht[20], and C-MPI[8].

It is important to point out that several key features of traditional DHTs are not necessary in HEC. Most HEC environments are batch oriented, which implies that a system that is configured at run time, generally has information about the compute and storage resources that will be available. This means that the amount of resources (e.g. number of nodes) would not change dynamically, and the only reason to decrease the allocation is either to handle failed nodes, or to terminate the allocation. Furthermore, nodes in HEC are generally reliable and have predictable uptime (nodes start on allocation, and nodes shut down on deallocation). This implies that node "churn" in HEC is virtually non-existent. HEC systems are generally locked down from the outside world, behind login nodes and firewalls, and although authentication and authorization is still needed, full communication encryption is wasteful for a large class of scientific computing applications that run on many HEC systems.

**Table -1 : DHT implementations.**

Name	Impl.	Routing Time	Persistence	Dynamic membership	Append
Cassandra 38	Java	log(N)	Yes	Yes	No
Memcached [20]	C	2	No	No	No
C-MPI [26]	C/MPI	log(N)	No	No	No
Dynamo [21]	Java	0 to log(N)	Yes	Yes	No
ZHT [14]	C++	0 to 2	Yes	Yes	Yes

There has been some uptake recently in using traditional DHTs in HEC, namely the C-MPI [8] project, in which the Kademia DHT has been implemented and shown to run well on 1000 nodes on a SiCortex machine. Another recent project using DHTs on a HEC is DataSpaces [7], which deploys a DHT on a Cray XT5 to coordinate in-memory data management for simulation workflows. Both C-MPI and DataSpaces fail to pay careful attention to latency by ensuring constant time operations. Furthermore, they do not decouple the metadata from the data management, causing potentially poor data locality. Amazon Dynamo [2] is arguably similar to the proposed ZHT, also claiming to be a zero-hop distributed hash-table. However, it's a non-opensource project, which makes its adoption not possible.

## 3. ZHT:AZERO HOP DISTRIBUTED HASHTABLE

We propose a new data-structure, named ZHT (Zero hop distributed Hash Table), which has been simplified and tuned for the specific requirements of High-End Computing (HEC). The ZHT features are described in this section.

**Hash Functions:** ZHT uses the SDBM hash function, due to its simple implementation, consistency across different data types (especially strings), and efficient performance.

**Membership Table:** The proposed hash functions map an arbitrary long string directly to an index value, which can then be used to efficiently retrieve the communication address (e.g. network name, IP address, MPI-rank) from a membership table.

**Failure Handling:** ZHT gracefully handles failures, by lazily tagging nodes that do not respond to requests repeatedly as failed (using exponential backoff). Once nodes are marked as down, they are assumed never to return until ZHT is restarted.

**Replication:** ZHT uses replication to ensure data stored persists during failures. New data created is pro-actively replicated to nodes in close proximity in the membership ID space. Ideally, once we support the network-aware topology of node ids (part of the future work section), the replication will consume the least amount of shared network resources by communicating only with neighbors in close proximity, and therefore improving scalability at extreme scales. We implemented replication on server side. When requests (insert and remove) are sent to servers, a thread will send the same request to corresponding replicas asynchronously. This will certainly introduce some overhead due to the sharing resource (CPU and network bandwidth). But our implementation don't introduce too much extra overhead when increase the number of replicas.

**Persistence:** ZHT is a distributed in-memory data-structure. We use a light-weight persistent hashtable with Kyotocabinet [6], which stores the ZHT state to persistent storage in real time.

### 3.1 Implementation

**Application Programming Interface:** The application programming interface (API) of ZHT is kept simple and follows similar interfaces for hash tables. The four operations ZHT will support are 1. bool *insert*(key, value); 2. value *retrieve*(key); 3. bool *remove*(key), and 4. bool *broadcast*(key, value). The insert operation will first hash the key to obtain the location of the node where the value will be inserted; existing values will be overwritten; replication occurs asynchronously between the hashed destination and its neighbors. The retrieve operation would return the value from ZHT if it existed. The remove operation would remove the value with the associated key. The broadcast operation would transmit the key/value pair over the edges of the spanning tree with the goal to distribute the key/value pair to all the caches.

**Implementation Details:** This work is aimed at HEC systems, so it was critical to identify a programming language supported unanimously across all current HEC systems; C/C++ was the best candidate from both a portability and performance perspective. FUSE and GridFTP is implemented in C, while the ZHT is in C++. GridFTP supports a variety of communication protocols, such as UDT and TCP; it supports GSS-API authentication of the control channel and data channel, and supports user-controlled levels of data integrity and/or confidentiality. ZHT used a proprietary binary UDP-based protocol; we are also investigating alternative communication protocols such as TCP, MPI, and BMI.

## 4. PERFORMANCE EVALUATION

In our implementation we completed an prototype of ZHT, which supports three of the four proposed operations, namely insert/lookup/remove. Prior to engaging in implementing ZHT, we explored the feasibility of using existing DHT implementations (Tapestry, Chord, Maidsafe-dht, and C-MPI), but came to the conclusion that they are all too heavy weight for HEC, were not feature complete, or had significant dependencies to resolve. The preliminary ZHT implementation was implemented in C++, and uses a proprietary binary TCP-based and UDP-based communication protocol.

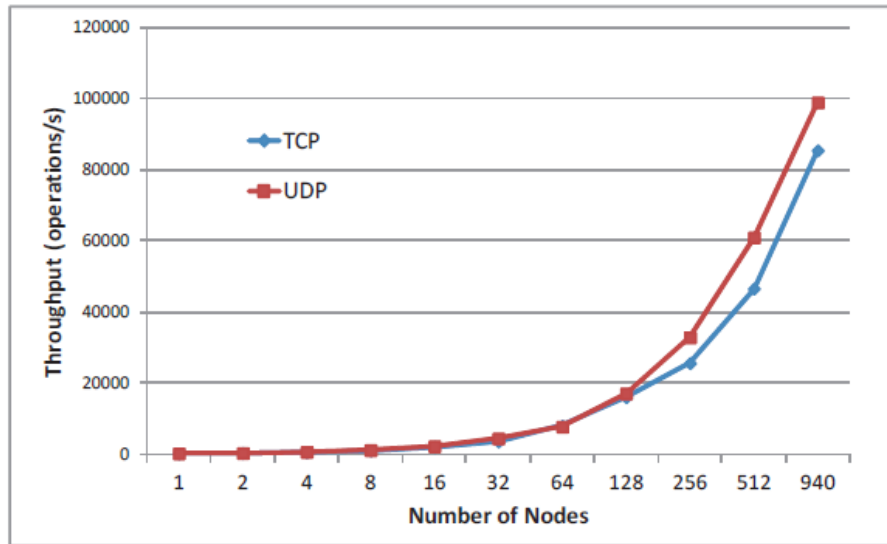


Chart -1: ZHT throughput on SiCortex

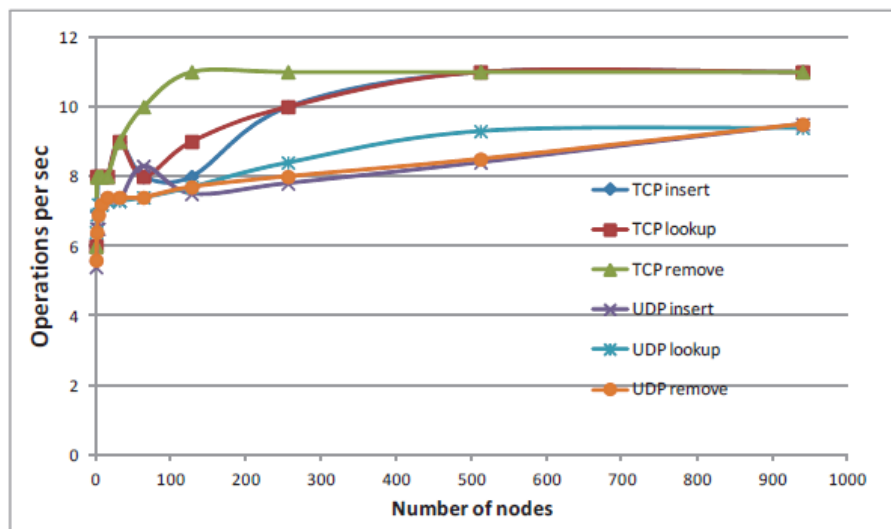


Chart -2: ZHT latency on SiCortex

We performed experiments (on the SiCortex SC5832 at ANL, 972 nodes with 5832 cores) to measure the overheads of insert/remove/lookup. Overheads were significantly lower than the DHTs that were investigated, with about 10ms per operation at modest scales of 5832-cores (as opposed to 30 ms for Chord and 1000 ms for Madsafe-DHT at 16-core scales). Figure 1 shows the throughput for ZHT for a range of scales from 1 node (6 cores) to 940 nodes (5640 cores). Each client (1 client per node) performs 100K random 132-byte key inserts, 100K lookups, and 100K removes. The throughputs in operations per second increases near linearly with scale, reaching 85K ops/sec at 940 nodes; the ideal throughput would have been 156K ops/sec. Also, as shown in the figure, the UDP protocol has better performance than TCP, which is near 100k ops/sec, and would might scale better in even large scale, nevertheless, based on the volatile propriety of UDP and bottleneck of network, UDP protocol in our implementation will lose performance when the network is saturated, after all, it does not hold a reliable way. Compared with the ideal linearly speed-up, ZHT hits more-than-half of the performance of it, it cannot be considered as a perfect work, however, it makes a giant progress in MTC.

Figure 2 shows the latency incurred by various operations on both TCP and UDP. TCP implementation keeps the latency within 11 ms while UDP does it with 9ms.

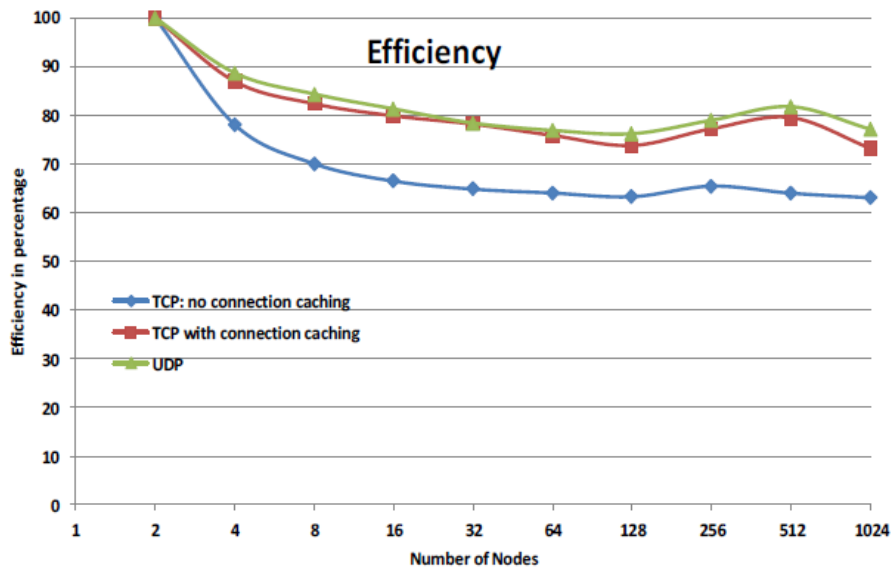


Chart -3: ZHT efficiency on SiCortex

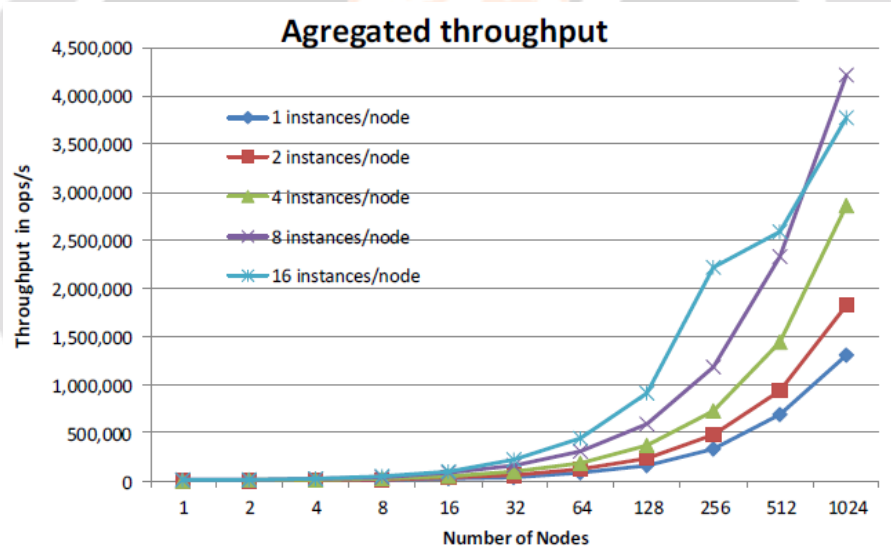


Chart -4: ZHT aggregated throughput on SiCortex

Additionally, we note that when the network saturates, TCP shows a better scalability and reliability. The performance differences among three basic operations (insert, lookup and remove) are very small and tend to be negligible. The time per operation starts at about 6ms at the smallest scales of 2 nodes (1 node scales is extremely fast because it has no network communication), and increases to 11ms to 12ms at the largest scale of 972 nodes. Since ZHT uses a direct 0-hop algorithm and that the majority of the overhead comes from network communication, it is not expected for the time per operation to increase significantly with larger scales.

### 5. CONCLUSIONS

The ideas in this position paper are transformative due to their departure from traditional HEC architectures and approaches, while proposing radical storage architecture changes based on distributed file systems to make exascale computing a reality. This paper addresses fundamental technical challenges that will become increasingly harder to



address with existing solutions due to a declining MTTF of future HEC systems. This work will open doors for novel research in programming paradigm shifts (e.g. Many-Task Computing [3]) needed as we approach exascale.

ZHT optimized for high-end computing systems was architected and implemented as a foundation in the development of fault-tolerant, high-performance, and scalable storage systems. We performed an extensive performance evaluation of ZHT. We have measured the performance of ZHT over TCP and UDP at scales up to 5400-cores on a SiCortex SC5832 supercomputer. ZHT is an open source project, available at [9].

Our work will benefit the 'Many-Task Computing' paradigm that bridges the gap between high-throughput computing and high-performance computing, generally producing both compute-intensive and data-intensive workloads, and has been shown to contain a large set of scientific computing applications from many domains.

Our main message is that by combining lessons learned from parallel file systems and distributed file systems, along with new advances in hardware (e.g. solid state memory), we can define a new storage architecture that is optimized for future high-end computing at exascale and has the potential to deliver a viable storage architecture for future extreme scale high-end computing. The position of this paper is revolutionary as it breaks the accepted practice of segregating storage resource from computational resources, and leveraging the abundance of processing power, bisection bandwidth, and local I/O commonly found in future high-end computing systems.

## 6. REFERENCES

- [1] V. Sarkar, et al. "ExaScale Software Study: Software Challenges in Extreme Scale Systems", ExaScale Computing Study, DARPA IPTO, 2009
- [2] F. Schmuck, R. Haskin, "GPFS: A Shared-Disk File System for Large Computing Clusters," FAST 2002
- [3] P. H. Carns, W. B. Ligon III, R. B. Ross, R. Thakur. "PVFS: A parallel file system for linux clusters", Proceedings of the 4th Annual Linux Showcase and Conference, 2000
- [4] P. Schwan. "Lustre: Building a file system for 1000-node clusters," Proc. of the 2003 Linux Symposium, 2003
- [5] I. Raicu, Z. Zhang, M. Wilde, I. Foster, P. Beckman, K. Iskra, B.Clifford. "Toward Loosely Coupled Programming on Petascale Systems," IEEE SC 2008
- [6] Z. Zhang, A. Espinosa, K. Iskra, I. Raicu, I. Foster, M. Wilde. "Design and Evaluation of a Collective I/O Model for Looselycoupled Petascale Programming", IEEE MTAGS08, 2008
- [7] S. Ghemawat, H. Gobioff, S.T. Leung. "The Google file system," 19th ACM SOSP, 2003
- [8] A. Bialecki, M. Cafarella, D. Cutting, O. O'Malley. "Hadoop: A Framework for Running Applications on Large Clusters Built of Commodity Hardware", 2005
- [9] H. Balakrishnan, M. F. Kaashoek, D. Karger, R. Morris, I. Stoica. "Looking up data in P2P systems", Communications of the ACM, 46(2):43-48, 2003
- [10] "Filesystem in Userspace", <http://fuse.sourceforge.net/>, 2011
- [11] W. Vogels, "Eventually consistent," ACM Queue, 2008.
- [12] I. Raicu, I. Foster, Y. Zhao. "Many-Task Computing for Grids and Supercomputers", IEEE MTAGS08, 2008
- [13] D. Zhao, C. Shou, X. Zhou, T. Li, Z. Zhang, I. Raicu, "FusionFS: a distributed \_lesystem for extreme scaledata intensive computing", Under MSST13 review, 2013.
- [14] T. Li, R. Verma, X. Duan, H. Jin, I. Raicu, "ZHT: Zero-Hop Distributed Hash Table for High-End Computing", ACM Performance Evaluation Review (PER), 2012
- [15] P. Maymounkov, D. Mazieres. "Kademlia: A Peer-to-peer Information System Based on the XOR Metric", In Proceedings of IPTPS, 2002
- [16] S. Ratnasamy, P. Francis, M. Handley, R. Karp, S. Schenker, "A scalable content-addressable network," in Proceedings of SIGCOMM, pp. 161-172, 2001
- [17] I. Stoica, R. Morris, D. Karger, M.F. Kaashoek, H. Balakrishnan, "Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications", ACM SIGCOMM, pp. 149-160, 2001
- [18] A. Rowstron and P. Druschel, "Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems," in Proceedings of Middleware, pp. 329-350, 2001
- [19] B.Y. Zhao, L. Huang, J. Stribling, S.C. Rhea, A.D. Joseph, and J.D. Kubiatowicz. "Tapestry: A Resilient Global-Scale Overlay for Service Deployment", IEEE Journal on Selected Areas in Communication, VOL. 22, NO. 1, 2004
- [20] B. Fitzpatrick. "Distributed caching with Memcached." Linux Journal, 2004(124):5, 2004
- [21] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, W. Vogels. "Dynamo: Amazon's Highly Available Key-Value Store." SIGOPS Operating Systems Review, 2007

- [22] H. Shen, C. Xu, and G. Chen. Cycloid: A Scalable Constant-Degree P2P Overlay Network. *Performance Evaluation*, 63(3):195-216, 2006
- [23] Ketama, <http://www.audioscrobbler.net/development/ketama/>, 2011
- [24] Riak, <https://wiki.basho.com/display/RIAK/Riak>, 2011
- [25] Maidsafe-DHT, <http://code.google.com/p/maidsafe-dht/>, 2011
- [26] J.M. Wozniak, B. Jacobs, R. Latham, S. Lang, S.W. Son, and R. Ross. "C-MPI: A DHT implementation for grid and HPC environments", Preprint ANL/MCS-P1746-0410, 2010
- [27] M. Wilde, I. Raicu, A. Espinosa, Z. Zhang, B. Clifford, M. Hategan, K. Iskra, P. Beckman, I. Foster. "Extreme scale scripting: Opportunities for large task-parallel applications on petascale computers", SciDAC09, 2009
- [28] C. Docan, M. Parashar, S. Klasky. "DataSpaces: An Interaction and Coordination Framework for Coupled Simulation Workflows", ACM HPDC 2010 786e 1

