

## COMPENDIUM ON LOAD RESTABILIZE FOR SCATTERED FILE SYSTEMS IN CLOUDS

**Darla Venkata Vara Prasad 1\*, N Anjaneyulu 2\***

1. II.M.Tech , Dept of CSE, AM Reddy Memorial College of Engineering & Technology, Petlurivaripalem.
2. Asst Prof, Dept. of CSE, AM Reddy Memorial College of Engineering & Technology, Petlurivaripalem.

**Abstract:** Distributed file systems is an important issue in DHTs. All the DHTs make some effort to load balance. Which keys get mapped; some machines get a larger portion of it. Thus, even if keys are numerous and random, some machines receive more than their fair share, by as much as a factor of  $n$  times the average. To cope with this problem, many DHTs use virtual nodes each real machine pretends to be several distinct machines, each participating independently in the DHT protocol. There are several provably efficient load balancing for distributed file's protocols for distributed data storage in P2P systems. More details and analysis can be found in this. Our algorithms are simple, and easy to implement in distributed files so an obvious next research step should be a practical evaluation of these schemes. Distributed does not easily generalize to more than one order.

**Keywords:** DHT, Centralize System, Load Im Balancing, Distributed System.

### 1. INTRODUCTION

The Distributed file systems an important issue in DHTs is load-balance the even distribution of items to nodes in the DHT. All DHTs make some effort to load balance; generally by randomizing the DHT address associated with each item with a "good enough" hash function and making each DHT node responsible for a balanced portion of the DHT address space. Chord is a prototypical example of this approach: its "random" hashing of nodes to a ring means that each node is responsible for only a small interval of the ring address space, while the random mapping of items means that only a limited number of items land in the (small) ring interval owned by any node. The cloud computing Current distributed hash tables do not evenly partition the address space into. Which keys get mapped; some machines get a larger portion of it. Thus, even if keys are numerous and random, some machines receive more than their fair share, by as much as a factor of  $n$  times the average. To cope with this problem, many DHTs use virtual nodes each real machine pretends to be several distinct machines, each participating independently in the DHT protocol. The machine's load is thus determined by summing over several virtual

nodes', creating a tight concentration of load near the average. As an example, the Chord DHT is based upon consistent hashing which requires virtual copies to be operated for every node. The node will occasionally check its inactive virtual nodes, and may migrate to one of them if the distribution of load in the system has changed. Since only one virtual node is active, the real node need not pay the original Chord protocol's multiplicative increase in space and bandwidth costs. Our solution to this problem therefore allows nodes to move to arbitrary addresses; with this freedom we show that we can load balance an arbitrary distribute Performance in a P2P System Our online load balancing algorithms are motivated by a new application domain for range partitioning peer-to-peer systems. P2P systems store a relation over a large and dynamic set of nodes, and support queries over this relation. Many current systems, known as Distributed Hash Tables (DHTs) use hash partitioning to ensure storage balance, and support point queries over the relation. There has been considerable recent interest in developing P2P systems that can support efficient range queries. For example, a P2P multi-player game might query for all

objects located in an area in a virtual 2-D space. In a P2P web cache, a node may request (pre-fetch) all pages with a specific URL prefix. It is well-known that hash partitioning is inefficient for answering such ad hoc range queries, motivating a search for new networks that allow range partitioning while still maintaining the storage balance offered by normal DHTs.

### Handling Dynamism in the Network

The network is a splits the range of  $N_h$  to take over half the load of  $N_h$ , using the NBRADJUST operation. After this split, there maybe NBRBALANCE violations between two pairs of neighbors and In response, ADJUSTLOAD is executed, first at node  $N_h$  and then at node  $N$ . It is easy to show (as in Lemma 3) that the resulting sequence of NBRADJUST operations repair all NBRBALANCE violations.

### Node Departure

While in the network, each node manages data for a particular range. When the node departs, the data is stored becomes unavailable to the rest of the peers. P2P networks reconcile this data loss in two ways: (a) Do nothing and let the "owners" of the data deal with its availability.

The owners will frequently poll the data to detect its loss and re-insert the data into the network.

Maintain replicas of each range across multiple nodes. A Skip Net DHT organizes peers and data objects according to their lexicographic addresses in the form of a variant of a probabilistic skip list. It supports logarithmic time range-based lookups and guarantees path locality. Mercury is more general than Skip Net since it supports range-based lookups on multiple-attributes. Our use of random sampling to estimate query selectivity constitutes a novel contribution towards implementing scalable multi-dimensional range queries. Load balancing is another important way in which Mercury from Skip Net. While Skip Net incorporates a constrained load-balancing mechanism, it is only useful when part of a data

name is hashed, in which case the part is inaccessible for performing a range query. This implies that Skip Net supports load-balancing or range queries not both.

## 111. SYSTEM MODEL

### (a) Data Popularity

Unfortunately, in many applications, a particular range of values may exhibit a much greater popularity in terms of database insertions or queries than other ranges. This would cause the node responsible for the popular range to become overloaded. One obvious solution is to determine some way to partition the ranges in proportion to their popularity.

As load patterns change, the system should also move nodes around as needed. We leverage our approximate histograms to help implement load-balancing in Mercury.

First, each node can use histograms to determine the average load existing in the system, and, hence, can determine if it is relatively heavily or lightly loaded. Second, the histograms contain information About which parts of the overlay are lightly loaded.

About which parts of the overlay are lightly loaded. (b)Load Balancing

We have shown how to balance the address space, but sometimes this is not enough. Some applications, such as those aiming to support range-searching operations, need to specify a particular, non-random mapping of items into the address space. In this section, we consider a dynamic protocol that aims to balance load for arbitrary item distributions. To do so, we must sacrifice the previous protocol's restriction of each node to a small number of virtual node locations. instead, each node is free to migrate anywhere. Our protocol is randomized, and relies on the underlying P2P routing framework only insofar as it has to be able to contact "random" nodes in the system (in the full paper we show that this can be done even when the node distribution is skewed by the load balancing protocol). The protocol is the following, to state the performance of the protocol, we need the concept of a half-life [LNBK02], which is the time it takes for half the

nodes or half the items in the system to arrive or depart.

### (c) DHT Implementation

The storage nodes are structured as a network based on distributed hash tables (DHTs), e.g., discovering a file chunk can simply refer to rapid key lookup in DHTs, given that a unique handle (or identifier) is assigned to each file chunk. DHTs enable nodes to selforganize and Repair while constantly offering lookup functionality in node dynamism, simplifying the system provision and management. The chunk servers in our proposal are organized as a DHT network. Typical DHTs guarantee that if a node leaves, then its locally hosted chunks are reliably migrated to its successor; if a node joins, then it allocates the chunks whose IDs immediately precede the joining ion of items, without expending much cost in maintaining the load balances.

Our scheme works through a kind of "work stealing" in which under loaded nodes migrate to portions of the address space occupied by too many items. The protocol is simple and practical, with all the complexity in its performance analysis. In this paper, we are interested in studying the load rebalancing problem in distributed file systems specialized for large-scale, dynamic and data-intensive clouds. Lastly, permitting nodes to choose arbitrary addresses in our item balancing protocol makes it easier for malicious nodes to disrupt the operation of the P2P network. It would be interesting to find counter measures for this problem.

The paper is organized as follows. Section II Related work, Section III. System Model Section IV Load balancing algorithm. Section V Distributed files system. Section VI. Performance Evaluation VII concludes.

## II. RELATED WORK

This attempt to load-balance can fail in two ways. First, the typical "random" partition of the address space among nodes is not completely balanced. Some nodes end up with

a larger portion of the addresses and thus receive a larger portion of the randomly distributed items. Second, some applications may preclude the randomization of data items' addresses. For example, to support range searching in a database application the items may need to be placed in a specific order, or even at specific addresses, on the ring. In such cases, we may find the items unevenly distributed in address space, meaning that balancing the address space among nodes is not adequate to balance the distribution of items among nodes. We give protocols to solve both of the load balancing challenges just described.

### Performance in a P2P System

Our online load balancing algorithms are motivated by a new application domain for range partitioning peer-to-peer systems. P2P systems store a relation over a large and dynamic set of nodes, and support queries over this relation. Many current systems, known as Distributed Hash Tables (DHTs) use hash partitioning to ensure storage balance, and support point queries over the relation. There has been considerable recent interest in developing P2P systems that can support efficient range queries. For example, a P2P multi-player game might query for all objects located in an area in a virtual 2-D space. In a P2P web cache, a node may request (pre-fetch) all pages with a specific URL prefix. It is well-known that hash partitioning is inefficient for answering such ad hoc range queries, motivating a search for new networks that allow range partitioning while still maintaining the storage balance offered by normal DHTs.

### Handling Dynamism in the Network

The network is a splits the range of  $N_h$  to take over half the load of  $N_h$ , using the NBRADJUST operation. After this split, there may be NBRBALANCE violations between two pairs of neighbors and In response, ADJUSTLOAD is executed, first at node  $N_h$  and then at node  $N$ . It is easy to show (as in Lemma 3) that the resulting sequence of NBRADJUST operations repair all NBRBALANCE violations.

## Node Departure

While in the network, each node manages data for a particular range. When the node departs, the data is stored becomes unavailable to the rest of the peers. P2P networks reconcile this data loss in two ways: (a) Do nothing and let the "owners" of the data deal with its availability.

The owners will frequently poll the data to detect its loss and re-insert the data into the network.

Maintain replicas of each range across multiple nodes. A Skip Net DHT organizes peers and data objects according to their lexicographic addresses in the form of a variant of a probabilistic skip list. It supports logarithmic time range-based lookups and guarantees path locality. Mercury is more general than Skip Net since it supports range-based lookups on multiple-attributes. Our use of random sampling to estimate query selectivity constitutes a novel contribution towards implementing scalable multi-dimensional range queries. Load balancing is another important way in which Mercury from Skip Net. While Skip Net incorporates a constrained load-balancing mechanism, it is only useful when part of a data name is hashed, in which case the part is inaccessible for performing a range query. This implies that Skip Net supports load-balancing or range queries not both.

### 11.1. SYSTEM MODEL (a) Data Popularity

Unfortunately, in many applications, a particular range of values may exhibit a much greater popularity in terms of database insertions or queries than other ranges. This would cause the node responsible for the popular range to become overloaded. One obvious solution is to determine some way to partition the ranges in proportion to their popularity.

As load patterns change, the system should also move nodes around as needed. We leverage our approximate histograms to help implement load-balancing in Mercury.

First, each node can use histograms to determine the average load existing in the system, and, hence, can determine if it is relatively heavily or lightly loaded. Second, the histograms contain information about which parts of the overlay are lightly loaded.

About which parts of the overlay are lightly loaded. (b) Load Balancing

We have shown how to balance the address space, but sometimes this is not enough. Some applications, such as those aiming to support range-searching operations, need to specify a particular, non-random mapping of items into the address space. In this section, we consider a dynamic protocol that aims to balance load for arbitrary item distributions. To do so, we must sacrifice the previous protocol's restriction of each node to a small number of virtual node locations. Instead, each node is free to migrate anywhere. Our protocol is randomized, and relies on the underlying P2P routing framework only insofar as it has to be able to contact "random" nodes in the system (in the full paper we show that this can be done even when the node distribution is skewed by the load balancing protocol). The protocol is the following, to state the performance of the protocol, we need the concept of a half-life [LNBK02], which is the time it takes for half the nodes or half the items in the system to arrive or depart.

### (c) DHT Implementation

The storage nodes are structured as a network based on distributed hash tables (DHTs), e.g., discovering a file chunk can simply refer to rapid key lookup in DHTs, given that a unique handle (or identifier) is assigned to each file chunk. DHTs enable nodes to self-organize and repair while constantly offering lookup functionality in node dynamism, simplifying the system provision and management. The chunk servers in our proposal are organized as a DHT network. Typical DHTs guarantee that if a node leaves, then its locally hosted chunks are reliably migrated to its successor; if a node joins, then it allocates the chunks whose IDs immediately precede the joining node from its successor to manage. Now we describe the application of this idea to DHTs. Let  $h_0$  be a

universally agreed hash function that maps peers onto the ring. Similarly, let  $h_1; h_2; \dots; h_d$  be a series of universally agreed hash functions mapping items onto the ring. To insert an item  $x$  using  $d$  hash functions, a peer calculates  $h_1(x), h_2(x), \dots, h_d(x)$ . Then,  $d$  lookups are executed in parallel to find the peers  $p_1; p_2, \dots, p_d$  responsible for these hash values, according to the mapping given by  $h_0$ .

#### (d) Chunk creation

A file is partitioned into a number of chunks allocated in distinct nodes so that Map Reduce Tasks can be performed in parallel over the nodes. The load of a node is typically proportional to the number of file chunks the node possesses. Because the files in a cloud can be arbitrarily created, deleted, and appended, and nodes can be upgraded, replaced and added in the file system, the file chunks are not distributed as uniformly as possible among the nodes. Our objective is to allocate the chunks of files as uniformly as possible among the nodes such that no node manages an excessive number of chunks.

#### (e) Replica Management

In distributed file systems (e.g., Google GFS and Hadoop HDFS), a constant number of replicas for each file chunk are maintained in distinct nodes to improve file availability with respect to node failures and departures. Our current load balancing algorithm does not treat replicas distinctly. It is unlikely that two or more replicas are placed in an identical node because of the random nature of our load rebalancing algorithm. More specifically, each under loaded node samples a number of nodes, each selected with a probability of  $1/n$ , to share their loads (where  $n$  is the total number of storage nodes).

### IV. LOAD BALANCING ALGORITHM

In our proposed algorithm, each chunk server node  $i$  first estimates whether it is under loaded (light) or overloaded (heavy) without global knowledge. A node is light if the number of chunks it hosts is smaller than the threshold.

Load statuses of a sample of randomly selected nodes.

Specifically, each node contacts a number of randomly selected nodes in the system and builds a vector denoted by  $V$ . A vector consists of entries, and each entry contains the ID, network address and load status of a randomly selected node. Fig. 1 shows the total number of messages generated by a load rebalancing algorithm, A large-scale distributed file system is in a

Load-balanced state: if each chunk server hosts no more than  $A_m$  chunks. In our proposed algorithm, each chunk server node  $i$  first estimates whether it is under loaded (light) or overloaded (heavy) without global knowledge. Let  $A_j$  be the load of node  $j$  to relieve  $j$ 's load. Node  $j$  may still remain as the heaviest node in the system after it has migrated its load to node  $i$ . In this case, the current least-loaded node, say node  $i$  departs and then rejoins the system as  $j$ 's successor. That is,  $i$  becomes node  $j+1$ , and  $j$ 's original successor  $i$  thus becomes node  $j + 2$ . Such a process repeats iteratively until  $j$  is no longer the heaviest. Then, the same process is executed to release the extra load on the next heaviest node in the system. This process repeats until all the heavy nodes in the system become light nodes.

Others: We will offer a rigorous performance analysis for the effect of varying  $nV$  in Appendix E. Specifically, we discuss the tradeoff between the value of  $nV$  and the movement cost. A larger  $nV$  introduces more overhead for message exchanges, but results in a smaller movement cost.

**Example 1:** In the setting above, the maximum load is at most  $\log \log n = \log d + 0$  with high probability. Our proof (not included for reasons of space) uses the layered induction technique from the seminal work of Because of the variance in the arc length associated with each peer; we must modify the proof to take this into account. The standard layered induction uses the fact that if there is  $k$  bins that have load at least  $k$ ,

**Example 2:** long distance links are constructed using the harmonic distribution on node link

distance. Value Link denotes the overlay when the harmonic distribution on value distance. Given the capacities of nodes (denoted by  $\{P_i, (32, \dots)\}$ ), we enhance the basic algorithm in Section III-B2 as follows: each node  $i$  approximates the ideal number of file chunks that it needs to host in a load balanced state as follows:

$$A_i = \gamma r_i^3$$

Note that the performance of the Value Link overlay is representative of the performance of a plain DHT under the absence of hashing and in the presence of load balancing algorithms which preserve value contiguity.

## V. DISTRIBUTED FILE SYSTEM

We have given several provably efficient load balancing for distributed file's protocols for distributed data storage in P2P systems. More details and analysis can be found in a thesis. Our algorithms are simple, and easy to implement in distributed files so an obvious next research step should be a practical evaluation of these schemes. In addition, several concrete open problems follow from our work. First, it might be possible to further improve the consistent hashing scheme as discussed at the end of our range search data structure. Distributed does not easily generalize to more than one order.

## VI. PERFORMANCE EVALUATION

We run a varying number of players. The players move through the world according to a random waypoint model, with a motion time chosen uniformly at random from seconds, a destination chosen uniformly at random, and a speed chosen uniformly at random from (0, 360) pixels per second. The size of the game world is scaled according to the number of players. The dimensions are  $640n \times 480n$ , where  $n$  is the number of players. All results are based on the average of 3 experiments, with each experiment lasting 60 seconds. The experiments include the bent of  $\log n$  sized LRU cache long pointers. the HDFS load balancer and our proposal. Our proposal clearly outperforms the HDFS load balancer.

When the name node is heavily loaded (i.e., small  $M$ 's), our proposal remarkably performs better than the HDFS load balancer. For example, if  $M = 1\%$ , the HDFS load balancer takes approximately 60 minutes to balance the loads of data nodes. By contrast, our proposal takes nearly 20 minutes in the case of  $M = 1\%$ . Specifically, unlike the HDFS load balancer, our proposal is independent of the load of the name node. In particular, approximating the unlimited scenario is expensive, and the use of  $\log_2 n$  virtual peers as proposed in introduces a large amount of topology maintenance traffic but does not provide a very close approximation. Finally, we observe that while we are illustrating the most powerful instantiation of virtual peers, we are comparing it to the weakest choice model further improvements are available to us just by increasing  $d$  to 4.

## VII. CONCLUSIONS

This paper gives the details of how to balance the loads of nodes and reduce the demanded movement cost as much as possible. In this the synthesis workloads stress test the load balancing algorithms by creating a few storage nodes that are heavily loaded.

## REFERENCES

1. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. E. Kaashoek, F. Dabek, and H. Balakrishnan, "Chord: a Scalable Peer-to-Peer Lookup Protocol for Internet Applications," IEEE/ACM Trans. Netw., vol. 11, no. 1, pp. 17.21, Feb. 2003.
2. A. Rowstron and P. Druschel, "Pastry: Scalable, Distributed Object Location and Routing for Large-Scale Peer-to-Peer Systems," LNCS 2218, pp. 161.172, Nov. 2001.
3. G DeCandia, D. Hastorun, M. Jampani, G Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels, "Dynamo: Amazon's Highly Available Key-value Store," in Proc. 21st ACM Symp. Operating Systems Principles (SOSP'07), Oct. 2007, pp. 205.220.

4. A. Rao, K. Lakshminarayanan, S. Surana, R. Karp, and I. Stoica, "Load Balancing in Structured P2P Systems," in Proc. 2nd Int'l Workshop Peerto-Peer Systems (IPTSEIf02), Feb. 2003, pp. 68.79.
5. D. Karger and M. Ruhl, "Simple Efficient Load Balancing Algorithms for Peer-to-Peer Systems,Dh in Proc. 16th ACM Symp. Parallel Algorithms and Architectures (SPAA'04), June 2004, pp. 36.43.
6. D. DeWitt, R. H. Gerber, G Graefe, M. L. Heytens, K. B. Kumar, and M.Muralikrishna. Gamma -a high performance dataflow database. In Proc. VLDB, 1986.
7. D.DeWitt and J. Gray. Parallel database systems: The future of high performance database processing. Communications of the ACM, 36(6), 1992.
8. H. Feelifl, M. Kitsuregawa, and B. C. Ooi. A fast convergence technique for online heat-balancing of btree indexed database over shared-nothing parallel systems. In Proc. DEXA, 2000.
9. P. Ganesan, B. Yang, and H. Garcia-Molina. One torus to rule them all: Multidimensional queries in p2p systems. In WebDB, 2004.
10. S. Ghandeharizadeh and D. J. DeWitt. A performance analysis of alternative multiattribute declustering strategies. In Proc. SIGMOD, 1992.
- N. J. A. Harvey, M. Jones, S. Saroiu, M. Theimer, and A. Wolman. Skipnet: A scalable overlay network with practical locality properties. In Proc. USITS, 2003.
12. D. R. Karger and M. Ruhl. Simple efficient load-balancing algorithms for peer-to-peer systems. In Proc. IPTPS, 2004.